A Concise Report of the 7th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text

Ali Hürriyetoğlu
Wageningen Food Safety Research, Netherlands
ali.hurriyetoglu@wur.nl

Surendrabikram Thapa
Virginia Tech
Blacksburg, USA
sbt@vt.edu

Gökçe Uludoğan
Bogazici University
Istanbul, Turkey
gokce.uludogan@bogazici.edu.tr

Somaiyeh Dehghan
Sabanci University
Istanbul, Turkey
somaiyeh.dehghan@sabanciuniv.edu

Hristo Tanev
European Commission, Joint Research Centre
Ispra, Italy
hristo.tanev@ec.europa.eu

Abstract

In this paper, we provide a brief overview of the 7th workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE) co-located with EACL 2024. This workshop consisted of regular papers, system description papers submitted by shared task participants, and overview papers of shared tasks held. This workshop series has been bringing together experts and enthusiasts from technical and social science fields, providing a platform for better understanding event information. This workshop not only advances text-based event extraction but also facilitates research in event extraction in multimodal settings.

1 Introduction

In today's digital era, the voluminous and readily available data on socio-political, economic, and environmental phenomena hold transformative potential for data-driven analysis in the social and human sciences (Hürriyetoğlu et al., 2021a; Chen et al., 2023). This wealth of data supports informed policymaking by providing comprehensive insights into complex issues ranging from political unrest and environmental disasters to global health and economic crises (Shu and Ye, 2023). The rising demand by governments, international organizations, and NGOs for high-quality, actionable information underscores the critical role of data in addressing challenges, aiding those impacted by crises, and enhancing citizen welfare in various domains (Hürriyetoğlu et al., 2021c).

Recent advancements underscore the importance of leveraging big data for social good, highlighting how data-driven approaches can revolutionize our understanding and intervention strategies in critical areas such as humanitarian aid, healthcare, and environmental protection. Evans et al. (2023) introduce a privacy-preserving data analysis system that offers both privacy guarantees for individuals and statistical validity for researchers, facilitating the use of sensitive data in social science research. Furthermore, crowd-sourced text analysis, as discussed by Benoit et al. (2016), represents a paradigm shift in how data is collected and analyzed, enabling rapid, flexible, and reproducible data production that matches expert-level accuracy.

The intersection of data science with policy analysis suggests a promising avenue for integrating data-driven insights into decision-making processes. Zhang et al. (2020) highlight the convergence of data science and policy analysis, indicating an emergent cross-disciplinary domain where data science accelerates and enriches policy research. This integration is crucial for developing effective policies that are informed by empirical data and grounded in a deep understanding of complex societal dynamics. Thus, it is evident that data science is instrumental in devising strategies to prevent or mitigate conflicts, deliver aid to affected populations, enhance the well-being of citizens, and safeguard them through a multitude of approaches. The public's resistance to COVID-19 measures during 2020-2022 (Prakash and Das, 2022; Fainstein et al., 2023) and the conflict between Russia and Ukraine (Bhandari et al., 2023; Thapa et al., 2022) serve as prime instances of the critical need to harness this data for real-world impact, highlighting the public's growing demand for

Proceedings of the 7th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2024), pages 248–255

timely information regarding mass gatherings and societal trends.

In this context, the workshop titled 'Challenges and Applications of Automated Extraction of Socio-political Events from Text' (CASE 2024) plays a pivotal role. Organized as part of EACL 2024, the seventh edition of this workshop marks the continuation of a workshop series that has been ongoing since its inception (Hürriyetoğlu et al., 2022, 2021b, 2020). This workshop aims to explore the advancements and hurdles in the automated extraction of socio-political events from textual data, offering a platform for discussing the latest research findings, innovative methodologies, and the future of automated text analysis in capturing and interpreting complex social phenomena. The workshop encompasses a range of activities, including presentations of accepted papers, shared tasks that challenge participants, and keynote speeches from leading experts in the field, providing valuable insights into the state of the art and fostering collaboration among researchers and practitioners. This paper is a brief overview giving insights into the wide range of activities at CASE 2024.

2 Accepted Papers

This year, seven papers were accepted out of twelve submissions. Below, we provide brief descriptions of accepted papers.

- Fellman et al. (2024) created a new dataset called FanConInfo of comic convention websites with cleaned HTML, rendered screenshots, and human annotations for event name, start date, end date, and location. The authors compared the performance of GPT-4 Vision, GPT-4 Text, and GPT-3.5 on the FanConInfo dataset. The findings revealed that the vision-based GPT-4 model outperformed the text-based versions, achieving an 85% accuracy in exact match, significantly higher than GPT-4 Text's 64% and GPT-3.5's 59%. This underscores the effectiveness of visual methods in extracting web data. The research highlighted the importance of integrating textual and visual data for improved web scraping and suggested multimodal comprehension as a key direction for future AI advancements.

- Loerakker et al. (2024) trained and evaluated several language models on Dutch tweets to analyze their ability to classify tweets that express discontent. They hypothesize that people expressing discontent are more likely to protest. The authors found that models specifically pretrained on Twitter data, like Bernice (DeLucia et al., 2022) and TwHIN-BERT (Zhang et al., 2022), substantially outperform other models in classifying discontent tweets. The results highlight the importance of selecting appropriate models trained on similar data to the task domain. Though discontent classification is nuanced, the authors show it can help filter relevant messages and identify possible protests if models optimized for low false positives are used.

- Bakker et al. (2024) proposed a novel pipeline to automatically extract timelines from decision letters of Dutch FOIA requests, using SpaCy[1] to extract dates and ChatGPT to extract and classify event phrases. The authors created a dataset of 100 manually annotated decision letters and showed that the pipeline achieved 94% date extraction accuracy. The key contribution is demonstrating how to leverage ChatGPT's few-shot learning capabilities to build an accurate timeline extractor for a low-resource domain using just a small annotated dataset, without needing extensive training. The proposed approach effectively extracts and classifies events into coherent timelines from decision letters.

- Tanev (2024) presented a new weakly supervised method for sentence-level event detection using linear prototype patterns and approximate pattern matching with BERT (Devlin et al., 2019) embeddings. The method involves creating a set of linear event detection patterns (e.g. 'disease outbreak', 'number people were infected') that serve as prototypes for events of interest. BERT's contextualized word representations are then utilized to find semantic similarities between these patterns and text fragments, allowing the identification of related event phrases with high lexical and syntactic variability. The approach was evaluated on detecting two event types – new disease cases and terrorist attacks– where it achieved promising F1 scores comparable to supervised systems. A key advantage of this BERT-based technique is that it combines

[1] spacy.io/

the interpretability of pattern-based methods with BERT's implicit semantic knowledge to effectively handle linguistic variations while avoiding extensive supervision.

- Olsen et al. (2024) presented a contrast of socio-political event datasets from political science and NLP fields, highlighting differences in abstraction, source accessibility, and temporal dynamics. The authors showed that political science datasets focus on abstract event representations, while NLP datasets offer precise textual annotations for event extraction. The discrepancies include the level of detail, availability of source texts, and dataset dynamism. Further, they showed that recent initiatives aim to integrate these approaches, enhancing the richness and applicability of event data, yet also caution against ethical and bias considerations in politically sensitive contexts.

- Dehghan and Yanikoglu (2024) evaluated ChatGPT's efficiency in identifying hate speech within Turkish tweets, contrasting its performance against BERTurk's (Schweter, 2020) supervised fine-tuning on the SIU2023-NST (Arın et al., 2023) dataset. Results demonstrate BERTurk's superior accuracy and lower mean squared error (MSE) in detecting hate speech and assessing its intensity over both zero-shot and few-shot learning approaches of ChatGPT. Despite ChatGPT's advanced capabilities, BERTurk's specificity for the task underlines the importance of model and prompt design, suggesting ChatGPT's potential as a supplementary tool for dataset annotation with careful prompt crafting.

- Uludoğan et al. (2024b) introduced TurkishHatePrintCorpus, a new dataset of over 6600 Turkish news articles annotated for hate speech against ethnic, national, or religious groups. They also developed a model called HateTargetBERT that combines BERT representations with linguistic features tailored to detecting hate speech against specific groups. Experiments demonstrate that HateTargetBERT performs comparably or better than BERT alone and substantially outperforms using just the linguistic features. The target-oriented linguistic features also enable explaining the model's predictions. By releasing the dataset, model code, and features, the authors provide an important new resource for studying hate speech and show that augmenting BERT with hate speech linguistic patterns for particular groups is an effective and interpretable approach to detecting such content in Turkish news.

3 Shared Tasks

This edition of the workshop featured two shared tasks. In addition to these two tasks, the workshop also welcomed submissions to previously organized shared tasks in earlier editions of CASE workshops.

3.1 Task 1: Climate Activism Stance and Hate Event Detection Shared Task

Realizing the important role of social media in global discussions on climate change, this shared task is built on the fact that there is a diversity of opinions, including different points of view and stances including instances of hate speech. By dissecting the varied perspectives on climate activism, the shared task aimed to advance capabilities of automated systems in processing and analyzing climate-related social media discourse, particularly on Twitter, now X. The task was divided into three subtasks: A) Hate Speech Detection, B) Targets of Hate Speech Detection, and C) Stance Detection, each addressing different facets of the discourse surrounding climate change activism on social media.

Subtask A (Hate Speech Detection) required participants to classify tweets as exhibiting hate speech or not. This subtask drew attention to the prevalence of aggressive and harmful language in online discussions on climate change, reflecting the intensity of emotions and opinions on the topic. Subtask B (Targets of Hate Speech Detection) delved deeper by identifying the hate speech targets within tweets, categorizing them into individuals, organizations, or communities. This nuanced approach aimed to understand the direction of hate and its potential impacts on targeted groups/individuals. Subtask C (Stance Detection) focused on detecting the stance expressed in tweets towards climate change, categorizing them as supportive, oppositional, or neutral. This subtask sheds light on the diverse viewpoints in the climate change debate, emphasizing the complexity of public opinion on this global issue.

Thapa et al. (2024a) provide an overview of the shared task along with brief detail on the methods used by participants. The shared task hosted on codalab[2] attracted over 100 teams, with 23 participants submitting results for Subtask A, 18 for Subtask B, and 19 for Subtask C, showcasing a wide range of methodologies and approaches. The participants' ranking was determined on the basis of the macro F1-score.

In Subtask A, the highest performance was achieved by the team CUET_Binary_Hackers (Farsi et al., 2024) with an impressive F1-score of 91.44%, indicating a high level of accuracy in detecting hate speech in tweets. Their approach, which included a variety of machine learning and deep learning models, emphasized the effectiveness of various advanced algorithms in processing and understanding the nuances of language used in social media discourse. The use of oversampling techniques and a range of feature extraction methods further highlighted the complexity of identifying hate speech and the need for advanced computational techniques in tackling this challenge.

Subtask B saw MasonPerplexity (Emran et al., 2024) securing the top position with an F1-score of 78.58%, demonstrating the challenge of accurately identifying the targets of hate speech in the context of climate activism tweets. Their approach involved the use of a weighted ensemble model, incorporating back translation techniques to address class imbalance. Their method highlighted the innovative strategies required to enhance model performance in the context of limited and imbalanced data. With BERTweet-Large (Nguyen et al., 2020), they were able to get the first position.

Finally, for Subtask C, ARC-NLP (Kaya et al., 2024) topped the leaderboard with the highest F1-score of 74.83%. They also used a modified version of BERTweet (Nguyen et al., 2020). Their method employed a short input tokenization length (96 tokens) and incorporated special tokens for tweet-specific elements. Overall, this subtask tried to aid a broader problem of stance detection which helps in understanding public opinion dynamics on pressing global issues like climate change.

The results and methodologies presented across all subtasks provide valuable insights into the state-of-the-art capabilities in processing and understanding social media discourse on climate change. The reliance on transformer-based models (Vaswani et al., 2017), Large-language models (Thapa et al., 2023b) and innovative data processing techniques across the subtasks reflects the advanced computational approaches required to tackle the complexities of natural language in social media. However, the subtasks also highlight ongoing challenges, such as dealing with data imbalance, understanding nuanced expressions of hate speech and stance, and the ethical considerations in automated content analysis.

3.2 Task 2: Hate Speech Detection in Turkish and Arabic Tweets

The HSD-2Lang Shared Task at CASE 2024 focused on detecting hate speech in Turkish and Arabic tweets, which is a significant problem on social media platforms. Divided into two subtasks, Subtask A aimed to identify hate speech in Turkish across various contexts, while Subtask B tackled hate speech detection in Arabic with limited data. Both subtasks were binary classification problems aimed at distinguishing hateful from non-hateful tweets. Uludoğan et al. (2024a) explain the overview of the shared task in detail.

Subtask A involved a dataset of Turkish tweets annotated for hate speech related to refugees, the Israel-Palestine conflict, and Anti-Greek discourse, with 9,140 tweets for training and 2,295 for testing. The objective was to develop a model capable of accurately identifying hate speech in these tweets, with performance evaluated using the F1 score across all topics. Subtask B presented a challenge in building hate speech models from a smaller and highly imbalanced dataset of Arabic tweets focusing on anti-refugee sentiment. The training set for Subtask B comprised 860 tweets, 82 of which were labeled as hateful, and the test set contained 522 tweets, 52 of them being hateful. Similarly to Subtask A, this subtask was also evaluated using the F1 score on the test data.

The shared task attracted 33 teams, with 10 submitting results for Subtask A and 5 for Subtask B. Participants employed various BERT-based models, highlighting the versatility and effectiveness of these models in processing and classifying social media content. The winning team in Subtask A, DetectiveReDASers (Qachfar et al., 2024), achieved an F1 score of 0.69645 using ConvBERTurk (Schweter, 2020) with a novel pooling strategy and cross-lingual data augmentation,

[2]The competition page can be found here: https://codalab.lisn.upsaclay.fr/competitions/16206.

demonstrating the impact of innovative approaches and ensemble methods on enhancing detection capabilities. The winning team of Subtask B, ReBERT (Yagci et al., 2024), achieved an F1 score of 0.68354 by fine-tuning AraBERTv0.2, originally developed by Antoun et al. (2020), which had been pretrained on around 60 million Arabic tweets. This updated AraBERT model incorporated emojis and common words previously excluded from its vocabulary and was applied directly without preprocessing. Notable performance differences were observed among systems using the same method, emphasizing the importance of hyperparameter tuning in improving model performance. The marginal performance gap between the two subtasks is intriguing, especially given that the Turkish dataset was three times larger than the Arabic set, highlighting the effectiveness of pretrained models. The findings of this shared task highlight the importance of model selection, tuning, and the impact of various preprocessing and hyper-parameter choices on detection capabilities, offering insights for future research in multilingual hate speech detection.

3.3 Extended Task: Multimodal Hate Speech Event Detection During Russia-Ukraine Crisis

This shared task was conducted for the first time in 2023 in CASE 2023 co-located at RANLP 2023. Following the massive interest in this task, this shared task saw an impressive number of impressions in 2024 as well. This task was structured into two subtasks aimed at detecting hate speech in text-embedded images and identifying the targets of such hate speech, with performance evaluated using the macro F1-score. The shared task, also hosted in codalab[3], attracted 73 registered participants and marked significant progress, achieving the best F1-scores of 87.27% and 80.05% in Subtask A and Subtask B, respectively. Thapa et al. (2024b) summarize the findings of different teams in this extended shared task.

In Subtask A, participants demonstrated remarkable achievements with CLTL (Wang and Markov, 2024) leading the pack by attaining an F1-score of 87.27%, setting a new benchmark for detecting hate speech in text-embedded images. This score beats the top-score by ARC-NLP (Sahin et al., 2023) from the same shared task in CASE 2023 (Thapa et al., 2023a; Hürriyetoğlu et al., 2023). CLTL (Wang and Markov, 2024) proposed a method for the Multimodal Hate Speech Event Detection Shared Task that combines separate text and image processing modules with a simple MLP and softmax layer, offering a flexible and efficient alternative to complex Large Vision Language Models (LVLMs). Their modular, MLP-based feature fusion approach not only set a competitive benchmark by achieving the first position but also demonstrated the importance of model simplicity and the potential for significant performance gains through fine-tuning.

Similarly, subtask B saw CLTL (Wang and Markov, 2024) once again achieving the top performance with an F1-score of 80.05%, showcasing the feasibility and effectiveness of their approach in identifying hate speech targets in a multimodal setting. Their approach yet again beats the highest leaderboard score from CASE 2023. The diverse methodologies and significant accomplishments reported in this shared task reflect the ongoing efforts to advance hate speech detection technologies. The results from both subtasks indicate a growing capability to not only detect hate speech in multimodal content but also understand its targets, contributing to safer digital environments.

4 Conclusion

In conclusion, the diverse range of papers and shared tasks presented at the workshop, from multimodal data analysis to fine-tuning language models for detecting hate speech and extracting event timelines, underscores the potential of automated text analysis in understanding complex socio-political phenomena. The workshop's emphasis on addressing real-world issues, such as understanding discourse related to climate change, online hate speech, and misinformation, through state-of-the-art computational techniques, not only sets new benchmarks for future research but also highlights the growing commitment within the community to leverage natural language processing for social good. In the coming years, the workshop will continue to advance the intersection of natural language processing and social good, promoting cutting-edge research and interdisciplinary collaboration to tackle complex socio-political challenges through automated text analysis.

[3]The competition page can be found here: https://codalab.lisn.upsaclay.fr/competitions/16203.

Acknowledgements

This work is supported by the European Research Council Politus Project (ID:101082050) and European Union's HORIZON projects EFRA (ID: 101093026) and ECO-Ready (ID: 101084201).

References

Wissam Antoun, Fady Baly, and Hazem Hajj. 2020. Arabert: Transformer-based model for arabic language understanding. In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 9–15.

İnanç Arın, Zeynep Işık, Seçilay Kutal, Somaiyeh Dehghan, Arzucan Özgür, and Berrin Yanikoğlu. 2023. Siu2023-nst-hate speech detection contest. In *2023 31st Signal Processing and Communications Applications Conference (SIU)*, pages 1–4. IEEE.

Femke Bakker, Ruben van Heusden, and Maarten Marx. 2024. Timeline extraction from decision letters using chatgpt. In *Proceedings of The Seventh Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Kenneth Benoit, Drew Conway, Benjamin E Lauderdale, Michael Laver, and Slava Mikhaylov. 2016. Crowd-sourced text analysis: Reproducible and agile production of political data. *American Political Science Review*, 110(2):278–295.

Aashish Bhandari, Siddhant B Shah, Surendrabikram Thapa, Usman Naseem, and Mehwish Nasim. 2023. Crisishatemm: Multimodal analysis of directed and undirected hate speech in text-embedded images from russia-ukraine conflict. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1993–2002.

Yan Chen, Kate Sherren, Michael Smit, and Kyung Young Lee. 2023. Using social media images as data in social science research. *New Media & Society*, 25(4):849–871.

Somaiyeh Dehghan and Berrin Yanikoglu. 2024. Evaluating chatgpt's ability to detect hate speech in turkish tweets. In *Proceedings of The Seventh Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Alexandra DeLucia, Shijie Wu, Aaron Mueller, Carlos Aguirre, Philip Resnik, and Mark Dredze. 2022. Bernice: a multilingual pre-trained encoder for twitter. In *Proceedings of the 2022 conference on empirical methods in natural language processing*, pages 6191–6205.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.

Al Nahian Bin Emran, Amrita Ganguly, Sadiya Sayara Chowdhury Puspo, Dhiman Goswami, and Md Nishat Raihan. 2024. MasonPerplexity at ClimateActivism 2024: Integrating Advanced Ensemble Techniques and Data Augmentation for Climate Activism Stance and Hate Event Identification. In *Proceedings of the 7th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Georgina Evans, Gary King, Margaret Schwenzfeier, and Abhradeep Thakurta. 2023. Statistically valid inferences from privacy-protected data. *American Political Science Review*, 117(4):1275–1290.

Susan S Fainstein, John Forester, Kevin Lujan Lee, Tiara Na'puti, Julian Agyeman, Nicholas J Stewart, Johannes Novy, Aysin Dedekorkut Howes, Paul Burton, Stefan Norgaard, et al. 2023. Resistance and response in planning: Edited by susan s. fainstein and john forester. *Planning Theory & Practice*, pages 1–39.

Salman Farsi, Asrarul Hoque Eusha, and Mohammad Shamsul Arefin. 2024. CUET_Binary_Hackers at ClimateActivism 2024: A Comprehensive Evaluation and Superior Performance of Transformer Models in Hate Speech Detection and Stance Classification for Climate Activism. In *Proceedings of the 7th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Evan Fellman, Jacob Tyo, and Zachary Lipton. 2024. The future of web data mining: Insights from multimodal and code-based extraction methods. In *Proceedings of The Seventh Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Ali Hürriyetoğlu, Osman Mutlu, Erdem Yörük, Farhana Ferdousi Liza, Ritesh Kumar, and Shyam Ratan. 2021a. Multilingual protest news detection - shared task 1, CASE 2021. In *Proceedings of the 4th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2021)*, pages 79–91, Online. Association for Computational Linguistics.

Ali Hürriyetoğlu, Hristo Tanev, Osman Mutlu, Surendrabikram Thapa, Fiona Anting Tan, and Erdem Yörük.

2023. Challenges and applications of automated extraction of socio-political events from text (CASE 2023): Workshop and shared task report. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text*, pages 167–175, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.

Ali Hürriyetoğlu, Hristo Tanev, Vanni Zavarella, Jakub Piskorski, Reyyan Yeniterzi, Osman Mutlu, Deniz Yuret, and Aline Villavicencio. 2021b. Challenges and applications of automated extraction of socio-political events from text (CASE 2021): Workshop and shared task report. In *Proceedings of the 4th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2021)*, pages 1–9, Online. Association for Computational Linguistics.

Ali Hürriyetoğlu, Hristo Tanev, Vanni Zavarella, Reyyan Yeniterzi, Osman Mutlu, and Erdem Yörük. 2022. Challenges and applications of automated extraction of socio-political events from text (CASE 2022): Workshop and shared task report. In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, pages 217–222, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Ali Hürriyetoğlu, Erdem Yörük, Osman Mutlu, Fırat Duruşan, Çağrı Yoltar, Deniz Yüret, and Burak Gürel. 2021c. Cross-context news corpus for protest event-related knowledge base construction. *Data Intelligence*, 3(2):308–335.

Ali Hürriyetoğlu, Vanni Zavarella, Hristo Tanev, Erdem Yörük, Ali Safaya, and Osman Mutlu. 2020. Automated extraction of socio-political events from news (AESPEN): Workshop and shared task report. In *Proceedings of the Workshop on Automated Extraction of Socio-political Events from News 2020*, pages 1–6, Marseille, France. European Language Resources Association (ELRA).

Ahmet Kaya, Oguzhan Ozcelik, and Cagri Toraman. 2024. ARC-NLP at ClimateActivism 2024: Stance and Hate Speech Detection by Generative and Encoder Models Optimized with Tweet-Specific Elements. In *Proceedings of the 7th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Meagan Loerakker, Laurens Müter, and Marijn Schraagen. 2024. Fine-tuning language models on dutch protest event tweets. In *Proceedings of The Seventh Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. BERTweet: A pre-trained language model for English tweets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 9–14, Online. Association for Computational Linguistics.

Helene Bøsei Olsen, Étienne Simon, Erik Velldal, and Lilja Øvrelid. 2024. Socio-political events of conflict and unrest: A survey of available datasets. In *Proceedings of The Seventh Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Ashish Viswanath Prakash and Saini Das. 2022. Explaining citizens' resistance to use digital contact tracing apps: A mixed-methods study. *International Journal of Information Management*, 63:102468.

Fatima Zahra Qachfar, Bryan E. Tuck, and Rakesh M. Verma. 2024. DetectiveReDASers at HSD-2Lang 2024: A new pooling strategy with cross-lingual augmentation and ensembling for hate speech detection in low-resource languages. In *Proceedings of The Seventh Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Umitcan Sahin, Izzet Emre Kucukkaya, Oguzhan Ozcelik, and Cagri Toraman. 2023. Arc-nlp at multimodal hate speech event detection 2023: Multimodal methods boosted by ensemble learning, syntactical and entity features. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text*, pages 71–78.

Stefan Schweter. 2020. BERTurk - BERT models for Turkish. https://zenodo.org/records/3770924.

Xiaoling Shu and Yiwan Ye. 2023. Knowledge discovery: Methods from data mining and machine learning. *Social Science Research*, 110:102817.

Hristo Tanev. 2024. Leveraging approximate pattern matching with bert for event detection. In *Proceedings of The Seventh Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Surendrabikram Thapa, Farhan Jafri, Ali Hürriyetoğlu, Francielle Vargas, Roy Ka-Wei Lee, and Usman Naseem. 2023a. Multimodal hate speech event detection-shared task 4, case 2023. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text*, pages 151–159.

Surendrabikram Thapa, Usman Naseem, and Mehwish Nasim. 2023b. From humans to machines: can chatgpt-like llms effectively replace human annotators in nlp tasks. In *Workshop Proceedings of the 17th International AAAI Conference on Web and Social Media*.

Surendrabikram Thapa, Kritesh Rauniyar, Farhan Ahmad Jafri, Shuvam Shiwakoti, Hariram Veeramani, Raghav Jain, Guneet Singh Kohli, Ali Hürriyetoğlu, and Usman Naseem. 2024a. Stance and hate event detection in tweets related to climate activism - shared task at case 2024. In *Proceedings of the 7th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*.

Surendrabikram Thapa, Kritesh Rauniyar, Farhan Ahmad Jafri, Hariram Veeramani, Raghav Jain, Sandesh Jain, Francielle Vargas, Ali Hürriyetoğlu, and Usman Naseem. 2024b. Extended multimodal hate speech event detection during russia-ukraine crisis - shared task at case 2024. In *Proceedings of the 7th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*.

Surendrabikram Thapa, Aditya Shah, Farhan Ahmad Jafri, Usman Naseem, and Imran Razzak. 2022. A multi-modal dataset for hate speech detection on social media: Case-study of russia-ukraine conflict. In *CASE 2022-5th Workshop on Challenges and Applications of Automated Extraction of Socio-Political Events from Text, Proceedings of the Workshop*. Association for Computational Linguistics.

Gökçe Uludoğan, Somaiyeh Dehghan, İnanç Arın, Elif Erol, Berrin Yanikoglu, and Arzucan Özgür. 2024a. Overview of the Hate Speech Detection in Turkish and Arabic Tweets (HSD-2Lang) Shared Task at CASE 2024. In *Proceedings of The Seventh Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Gökçe Uludoğan, Atıf Emre Yüksel, Ümit Can Tunçer, Burak İşık, Yasemin Korkmaz, Didar Akar, and Arzucan Özgür. 2024b. Detecting Hate Speech in Turkish Print Media: A corpus and a hybrid approach with target-oriented linguistic knowledge. In *Proceedings of The Seventh Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Yeshan Wang and Ilia Markov. 2024. CLTL@Multimodal Hate Speech Event Detection 2024: The Winning Approach to Detecting Multimodal Hate Speech and Its Targets. In *Proceedings of the 7th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Utku Ugur Yagci, Ahmet Emirhan Kolcak, and Egemen Iscan. 2024. Rebert at HSD-2Lang 2024: Finetuning BERT with AdamW for hate speech detection in Arabic and Turkish. In *Proceedings of The Seventh Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Malta. Association for Computational Linguistics.

Xinyang Zhang, Yury Malkov, Omar Florez, Serim Park, Brian McWilliams, Jiawei Han, and Ahmed El-Kishky. 2022. Twhin-bert: a socially-enriched pre-trained language model for multilingual tweet representations. *arXiv preprint arXiv:2209.07562*.

Yi Zhang, Alan L Porter, Scott Cunningham, Denise Chiavetta, and Nils Newman. 2020. Parallel or intersecting lines? intelligent bibliometrics for investigating the involvement of data science in policy analysis. *IEEE Transactions on Engineering Management*, 68(5):1259–1271.