

Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2023): Workshop and Shared Task Report

Ali Hürriyetöglü

KNAW Humanities Cluster DHLab
name.surname@dh.huc.knaw.nl

Hristo Tanev

European Commission
hristo.tanev@ec.europa.eu

Osman Mutlu

Koc University
omutlu@ku.edu.tr

Surendrabikram Thapa

Virginia Tech
surendrabikram@vt.edu

Fiona Anting Tan

National University of Singapore
tan.f@u.nus.edu

Erdem Yörük

Koc University
eryoruk@ku.edu.tr

Abstract

We provide a summary of the sixth edition of the CASE workshop that is held in the scope of RANLP 2023. The workshop consists of regular papers, three keynotes, working papers of shared task participants, and shared task overview papers. This workshop series has been bringing together all aspects of event information collection across technical and social science fields. In addition to contributing to the progress in text based event extraction, the workshop provides a space for the organization of a multimodal event information collection task.

1 Introduction

Nowadays, the unprecedented quantity of easily accessible data on social, political, and economic processes offers ground-breaking potential in guiding data-driven analysis in social and human sciences and in driving informed policy-making processes. Governments, multilateral organizations, and local and global NGOs present an increasing demand for high-quality information about a wide variety of events ranging from political violence, environmental catastrophes, and conflict, to international economic and health crises (Coleman et al., 2014; Della Porta and Diani, 2015) to prevent or resolve conflicts, provide relief for those that are afflicted, or improve the lives of and protect citizens in a variety of ways. The citizen actions against the COVID measures in the period 2020-2022 and the war between Russia and Ukraine are only two examples where we must understand, analyze, and improve the real-life situations using such data. Finally, these efforts respond to “growing public interest in up-to-date information on crowds” as well.¹

The workshop Challenges and Applications of Automated Extraction of Socio-political Events

¹<https://sites.google.com/view/crowdcountingconsortium/faqs>

from Text (CASE 2023) is held in the scope of the conference Recent Advances in Natural Language Processing (RANLP). CASE 2023 is the sixth edition of a workshop series (Hürriyetöglü et al., 2022b, 2021b; Hürriyetöglü et al., 2020).

We provide brief notes about the accepted papers, shared tasks, and keynote speeches in the following sections.

2 Accepted papers

This year, all seven submissions were accepted by the program committee. A quick summary of these papers are provided below.

- Osorio and Vásquez (2023) collect and annotate a dataset in 3 different granularity: whether a document is related to criminal activities, whether each sentence is related to lethal behaviors, and each sentence’s membership to 11 predefined categories of events. Following these granularities, the authors design three binary classification tasks and apply numerous non-neural and neural models to the annotated dataset; they observe good performance and provide analysis for data slices with lower performance.
- Tanev and Longueville (2023) discern “main” location, where the event in question occurred, versus “secondary” locations, that provide extra context such as the origin of the protestors and the first event location in a series of protests. They accomplish this by training a BERT model on news articles annotated with the main event location (CASE 2021 shared task news dataset). The secondary locations are all other mentioned locations. They compared their results to that of two SVM models and a baseline that assumes only the first sentence contains the main location. Their model

outperformed all other systems with an F1 of 0.8 and accuracy 0.73.

- [Delucia et al. \(2023\)](#) proposes an extension of the Multiple-instance learning (MIL) framework to better handle a common problem in computational social science: given noisy reports from text sources, how can researchers identify if a bag of reports (here, tweets within a country-day) report a true event. The authors show that MIL improves civil unrest detection over methods based on simple aggregation. The experiments conducted on hyperparameters (key instance ratio n and instance supervision b) show an improvement from MIL- n compared to vanilla MIL and other variants by model selection.
- [Mbouadeu et al. \(2023\)](#) This paper studies the news headlines event linking task that given a title of news (typically a sentence), maps it to an event concept from a knowledge graph. The challenge is how to compare different (zero-shot) models' performance. They propose a benchmark for the evaluation and compare multiple models. By comparing three families of approaches (a) similarity based on rule or embeddings, (b) off-the-shelf entity linking tools, and (c) prompting Large Language Models, the authors show that the approach (c) has the best performance even though different approaches could be complementary to each other.
- [Slavcheva et al. \(2023\)](#) presents a semantic model to structure protest event ontology, and provides some general description of the practical work with the Bulgarian data. The paper presents both the modelling framework and the implementation. The model is a practical application of the Unified Eventy Representation (UER) formalism, which is based on the Unified Modeling Language (UML), whose four-layer architecture (i.e., user objects, model, metamodel, and meta-metamodel) provides flexible means for building the semantic representations of the language units along a scale of generality and specificity.
- [Tuparova et al. \(2023\)](#) proposes a method that extends the detection capability of existing event detection models to new event

types. The authors have an experimental setup for few-shot learning when there is a limited training sources. Moreover, they provide several analyses on the experimental results. The main strength is in using low resource (single GPU) in order to fine tune the model in a reasonable time by providing a small set of samples by leveraging the transfer learning feature of the pre-trained model. The performance on the detection of the already known events tends to improve as well.

- [Mutlu and Hürriyetoğlu \(2023\)](#) The paper propose a solution to address the issue of data scarcity in closed-domain event extraction. The proposed solution leverages on the use of a side-product of data annotation campaigns that are the data containing no annotation, by considering this information as a discriminant that improves the extraction performance. The authors propose a multi-task model where they leverage additional data present after the token annotation process. Experiments are well conducted, by showing the efficacy of the method using different gradually decreasing dataset dimensions.

3 Shared tasks

3.1 Task 1: Multilingual Protest News Detection

The performance of an automated system depends on the target event type as it may be broad or potentially the event trigger(s) can be ambiguous. The context of the trigger occurrence needs to be considered as well. For instance, depending on the context, the 'protest' event type may or may not be synonymous with 'demonstration'. Moreover, the hypothetical cases such as future protest plans may need to be excluded from the results. Finally, the relevance of a protest depends on the actors, since only citizen-led events are in the scope of contentious political events. This challenge becomes even harder in a cross-lingual and zero-shot setting where training data are not available in new languages. We tackle the task in four steps and hope state-of-the-art approaches will yield optimal results.

This shared task was announced as a re-run of the same tasks from CASE 2021 ([Hürriyetoğlu et al., 2021a](#)) and CASE 2022 ([Hürriyetoğlu et al., 2022a](#)). Although it attracted some interest, we

did not receive any task description papers for this edition.

3.2 Task 2: Automatically Replicating Manually Created Event Datasets

The purpose of Task 2 is to test the abilities of the event extraction systems to map events on the World map, by finding the locations where they have taken place. This year the subtitle of the task was "Detecting and Geocoding Battle Events from Social Media Messages on the Russo-Ukrainian War": The purpose of the task was to detect armed clash events in Russian and Ukrainian Telegram messages and to geocode them, i.e., find the location of the battles at the level of a populated place (Tanev et al., 2023a).

Until recently, event geocoding has been considered a topic, covered by the works in the area of Geoinformation systems. Only in the last years the NLP community started to consider ML algorithms for geocoding (Halterman, 2023), also in the context of event detection (Tanev et al., 2023b). In this context, we can consider our shared task as an evaluation exercise for such event geocoding systems.

The evaluation of the systems participating at shared task 2 relies on an original evaluation methodology which compares the battle coordinates, found by the systems with the locations of such events from a Gold standard data set.

As a Gold standard this year we used a selected subset of the ACLED event database (Raleigh et al., 2010), covering the first six months start of the Russo - Ukrainian conflict, namely 24 February 2022 - 24 August 2022 and considering only the events of type *battle*.

We provided the participants with English language text messages from Telegram channels originating from Russia and Ukraine. We gathered nearly 326K original English Telegram Messages from six Telegram channels.

The Telegram data is available from a public Github repository ².

Two systems participated in this year's evaluation: The top ranked system relied on a combination of a XLMRoberta (Liu et al., 2019) classifier, trained on ACLED, and a rule based geocoder using the JRC NEROne named entity recognition system (Jacquet et al., 2019). The second ranked

system used the NEXUS (Tanev et al., 2008) keyword based event classifier and Mordecai3 geoparser (Halterman, 2023).

The first system, based on XLMRoberta, achieved a moderate level of correlation with the ACLED dataset, which is a good result, considering the possibly low coverage of this year Telegram dataset, regarding Russo Ukrainian war.

We hypothesize that there is a low coverage of the Telegram dataset this year, since in the 2021 issue of this task (Giorgi et al., 2021), participating systems achieved much higher correlation with the ACLED Gold standard, including the NEXUS system. The low correlation can be explained with the low coverage of the war, where battles are not always reported in the media, especially in Telegram. There is a lot of imprecise information on the social media and in contrast ACLED gold standard relies on a wide range of verified sources, including radio and TV, and manually curates the data.

Our conclusion from this year database replication task was that social media is not the best source of information about armed conflicts.

3.3 Task 3: Event Causality Identification

Causality is a core cognitive concept and appears in many natural language processing (NLP) works that aim to tackle inference and understanding. We are interested in studying event causality in the news and, therefore, introduce the Causal News Corpus (Tan et al., 2022b). The dataset comprises of 3,767 event sentences extracted from protest event news, that have been annotated with sequence labels on whether it contains causal relations or not. Subsequently, causal sentences are also annotated with Cause, Effect and Signal spans. Two corresponding subtasks were involved in our shared task: In Subtask 1, participants were challenged to predict if a sentence contains a causal relation or not. In Subtask 2, participants were challenged to identify the Cause, Effect, and Signal spans given an input causal sentence. We hope that our shared task promotes research on the topic of detection and extraction of causal events in news.

This year's competition is the second iteration of the shared task, first introduced in 2022 (Tan et al., 2022a), and uses the latest version of the Causal News Corpus (CNC-V2), also known as RECESS (Tan et al., 2023a). As compared to V1 comprising of 160 sentences and 183 relations, the V2 contains 1,981 sentences and 2,754 causal relations for

²<https://github.com/htanev/RussoUkrainianWarTelegram>

Subtask 2. Annotations were also revised for some examples across both subtasks.

Tan et al. (2023b) provides an overview of the work of the ten teams that submitted their results to our competition and the six system description papers that were received. The top F1 score for Subtask 1 was 84.66% by Team DeepBlueAI, who did not submit a description paper. Team InterosML (Patel, 2023) scored a similar high score of 84.36%, and used a two step approach: first pre-training a baseline RoBERTa model with supervised contrastive loss, then fine-tuning the model on Subtask 1 itself. The top F1 score for Subtask 2 was 72.79% by Team BoschAI (Schrader et al., 2023), who used a sequence tagging approach to fine-tune BERT-large and RoBERTa-large, and adapted the target labels to allow prediction of up to three different causal relations per sentence.

3.4 Task 4: Multimodal Hate Speech Event Detection

Hate speech detection is one of the most important aspects of event identification during political events like invasions (Thapa et al., 2022). In the case of hate speech event detection, the event is the occurrence of hate speech, the entity is the target of the hate speech, and the relationship is the connection between the two. Since multimodal content is widely prevalent across the internet, the detection of hate speech in text-embedded images is very important.

Given a text-embedded image, task 4 aims to automatically identify the hate speech and its targets³. This task had two subtasks (Thapa et al., 2023). In subtask 1, participants were given a dataset of text-embedded images and the participants had to classify whether the given image contained hate speech or not. It was tasked as a binary classification problem of classifying hate speech and non-hate speech. Similarly, in subtask 2, the participants were given a dataset of hateful text-embedded images where they had to classify what the targets of hate speech were. This subtask was posed as a multi-class classification problem where targets were individual, community, and organization. The dataset curated by Bhandari et al. (2023) was used in this task. More than 50 participants registered for the competition.

Thapa et al. (2023) presents the overview of

³<https://codalab.lisn.upsaclay.fr/competitions/13087>

the performance of 13 teams who submitted their scores in subtask 1 and 10 teams who submitted their scores in subtask 2. The ranking was done on the basis of the macro F1-score. The competition saw a wide range of methodologies ranging from traditional machine learning models to powerful transformer architectures. The first team ARC-NLP (Sahin et al., 2023) proposed an ensemble of multilayer perceptions (for representations from textual and visual encoders) and various boosting algorithms (using syntactical and Bag-of-words representations) for subtask 1. The team was able to score an F1-score of 85.65%. Similarly, for subtask 2, they used Named Entity Recognition (NER) features along with CLIP representations. An ensemble approach similar to subtask 1 was able to give them the first position with an F1-score of 76.34%.

Similarly, many teams used transformer-based approaches. Out of the submitted papers for subtask 1, IIC_Team (Singh et al., 2023) ranked at rank 3 (F1-score 84.63%), Ometeotl (Armenta-Segura et al., 2023) at rank 6 (F1-score of 80.97%), and VerbaVisor (Esackimuthu and Balasundaram, 2023) at rank 8 (F1-score of 78.21%) were able to get the best performances with XLM-Roberta-base, BertForSequence classification, and ALBERT models respectively. All of them used the text extracted from the given dataset of text-embedded images using Google Vision API. In subtask 2, IIC_Team, VerbaVisor, and Ometeotl were able to get the rank of 3 with an F1-score of 69.73%, rank 4 with an F1-score of 68.05% and rank 7 with F1-score of 56.88% respectively with same models used in subtask 1.

Often, the visual information is also necessary. The first team utilized both textual and visual information effectively to get a high F1-score. Two teams viz. CSECU-DSG and LexicalSquad leveraged both textual and visual information. CSECU-DSG (Aziz et al., 2023) used a combination of BERT and vision transformers (ViT) (Dosovitskiy et al., 2020) to leverage textual and visual information respectively. They were able to get an F1-score of 82.48% and 65.30% in subtask 1 and subtask 2 respectively. They were placed at the fifth rank in both subtasks. Similarly, LexicalSquad (Kashif et al., 2023) participated only in subtask 1 where they used XLNet and BERT for textual features and Inception-V3 for visual features. With this combined representation, they were able to get an

F1-score of 74.96%. This ranked them at the tenth position in the leaderboard.

Traditional machine learning algorithms were also used by some teams where they performed decently well. SSN-NLP-ACE and ML_Ensemblers used various traditional machine learning approaches. SSN-NLP-ACE (K et al., 2023) used TF-IDF features with SVM (with RBF kernel) to get an F1-score of 78.80% in subtask 1 ranking them in the seventh position in the leaderboard. They used TF-IDF features with logistic regression for subtask 2 which ranked them at eighth position with an F1-score of 52.58%. Similarly, ML_Ensemblers used a variety of algorithms like Naive Bayes, KNN, SVM, and Decision Trees out of which Naive Bayes performed the best in both subtasks with an F1-score of 42.94% and 43.32% in subtask 1 and subtask 2 respectively. They were able to secure the rank of 13 and 9 in subtask 1 and subtask 2 respectively.

4 Keynotes

Three scholars delivered three keynote speeches that are summarized below.

4.1 Using Automated Text Processing to Understand Social Movements and Human Behaviour

Erdem Yörük’s keynote will describe two large-scale ERC-funded projects that employ computational social science methods to extract data on protests and public opinion. The first is the Global Contentious Politics Dataset (GLOCON) Project.⁴ is the first automated comparative protest event database on emerging markets using local news sources (Duruşan et al., 2022). The countries included in the GLOCON dataset are India, South Africa, Argentina, Brazil and Turkey. Glocon has been created by using natural language processing, and machine learning in order to extract protest data from online news sources. The project develops fully automated tools for document classification, sentence classification, and detailed protest event information extraction that performs in a multi-source, multi-context protest event setting with consistent performances of recall and precision for each country context. GLOCON counts the number of events such as strikes, rallies, boycotts, protests, riots, and demonstrations,

⁴<https://glocon.ku.edu.tr/>

i.e. the “repertoire of contention,” and operationalizes protest events by various social groups. The project has developed a novel bottom-up methodology that is based on a random sampling of news archives, as opposed to keyword filtering. The high-quality Gold standard corpus is designed in a way that can accommodate context variability from the outset as it is compiled randomly from a variety of news sources from different countries (Hürriyetoğlu et al., 2021; Yörük et al., 2021). The second one, Politus Project, aims at scaling up traditional survey polls for public opinion research with AI-based social data analytics. Politus develops an AI-based innovation that combines quantitative and computational methods to create a data platform that delivers representative, valid, instant, real-time, multi-country, and multi-language panel data on key political and social trends. The project will collect content information from Twitter and process it with AI tools to generate a large set of indicators on political and social trends through its data platform. The deep learning models and NLP tools will be designed from the ground up as language-independent and generalizable systems. The platform will deliver geolocated hourly panel data on demography, ideology, topics, values, and beliefs, behavior, sentiment, emotion, attitudes, and stance of users aggregated at the district level. In this keynote, Dr. Yörük will describe the general methodology of the projects, including data collection, data analysis, and their approach for representativeness, which is based on multilevel regression with post-stratification.

4.2 Bulgarian Event Corpus for the Construction of a Bulgaria-centric Knowledge Graph

The Bulgarian Event Corpus is being constructed within the CLaDA-BG (Bulgarian National Interdisciplinary Research E-Infrastructure for Bulgarian Language and Cultural Heritage Resources and Technologies. In the spirit of European CLARIN and DARIAH) we aim to support researchers in Humanities and Social Sciences (H&SS) to access the necessary datasets for their research. The different types of objects of study, representation and search are integrated on the basis of common metadata and content categories. The approach for interlinking of the datasets is called contextualization. The implementation of contextualization in CLaDA-BG will utilize a common Bulgaria-centered knowl-

edge graph - BGKG. The knowledge facts within BGKG are constructed around events of different types. Thus, construction of BGKG requires a set of appropriate language resources for training of Bulgarian language pipeline for extraction of events from text documents. A key element within these language resources is the Bulgarian Event Corpus. In the talk I will present the design of the annotation schema, the annotation process, relation to ontologies and RDF representation. We have started with the CIDOC-CRM ontology for the construction of the annotation schema. This ontology provides a good conceptualization of events motivated by the domain of museums which is appropriate for our goals. During the design of the annotation schema, we extended the ontology with new events depending on the content of the corpus. The documents to be annotated were selected from scientific and popular publications of the partners within CLaDA-BG and articles from Bulgarian Wikipedia. The annotation is done on several layers: Named Entities, Events, Roles, Linking, terms and keywords.

4.3 With a little help from NLP: My Language Technology applications with impact on society

Ruslan Mitkov will present original methodologies developed by the speaker, underpinning implemented Language Technology tools which are already having an impact on the following areas of society: e-learning, translation and interpreting and care for people with language disabilities.

The first part of the presentation will introduce an original methodology and a tool for generating multiple-choice tests from electronic textbooks. The application draws on a variety of Natural Language Processing (NLP) techniques which include term extraction, semantic computing and sentence transformation. The presentation will include an evaluation of the tool which demonstrates that generation of multiple-choice tests items with the help of this tool is almost four times faster than manual construction and the quality of the test items is not compromised. This application benefits e-learning users (both teachers and students) and is an example of how NLP can have a positive societal impact, in which the speaker passionately believes. The latest version of the system based on deep learning techniques will also be briefly introduced.

The talk will go on to discuss two other original

recent projects which are also related to the application of NLP beyond academia. First, a project, whose objective is to develop next-generation translation memory tools for translators and, in the near future, for interpreters, will be briefly presented. Finally, a project will be outlined which focuses on helping users with autism to read and better understand texts. The speaker will put forward ideas as to what we can do next.

The presentation will finish with a brief outline of the latest (and forthcoming) research topics (to be) which the speaker plans to pursue and his vision on the future NLP applications. In particular, he will share his views as to how NLP will develop and what should be done for NLP to be more successful, more inclusive and more ethical.

5 Conclusion

Many aspects of event information modeling and collection are reported in the scope of CASE 2023. Hosting a shared task that is on multimodal problem and having submissions about languages other than English (e.g., Bulgarian) are distinguishing aspects of this edition.

References

- Jesus Armenta-Segura, César Jesús Núñez-Prado, Grigori Olegovich Sidorov, Alexander Gelbukh, and Rodrigo Francisco Román-Godínez. 2023. Ometeotl@multimodal hate speech event detection 2023: Hate speech and text-image correlation detection in real life memes using pre-trained bert models over text. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Abdul Aziz, Md. Akram Hossain, and Abu Nowshed Chy. 2023. Csecu-dsg@multimodal hate speech event detection 2023: Transformer-based multimodal hierarchical fusion model for multimodal hate speech detection. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- 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 Workshops*, pages 1993–2002.

- Peter T Coleman, Morton Deutsch, and Eric C Marcus. 2014. *The handbook of conflict resolution: Theory and practice*. John Wiley & Sons.
- Donatella Della Porta and Mario Diani. 2015. *The Oxford handbook of social movements*. Oxford University Press.
- Alexandra Delucia, Mark Dredze, and Anna L. Buczak. 2023. A multi-instance learning approach to civil unrest event detection on twitter. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*.
- Firat Duruşan, Ali Hürriyetoğlu, Erdem Yörük, Osman Mutlu, Çağrı Yoltar, Burak Gürel, and Alvaro Comin. 2022. [Global contentious politics database \(glocon\) annotation manuals](#).
- Sarika Esackimuthu and Prabavathy Balasundaram. 2023. Verbavisor@multimodal hate speech event detection 2023: Hate speech detection using transformer model. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Salvatore Giorgi, Vanni Zavarella, Hristo Tanev, Nicolas Stefanovitch, Sy Hwang, Hansi Hettiarachchi, Tharindu Ranasinghe, Vivek Kalyan, Paul Tan, Shaun Tan, et al. 2021. Discovering black lives matter events in the united states: Shared task 3, case 2021. In *Proceedings of the 4th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2021)*, pages 218–227.
- Andrew Halterman. 2023. Mordecai3: A neural geoparser and event geolocator. *working paper*.
- Ali Hürriyetoğlu, Osman Mutlu, Firat Duruşan, Onur Uca, Alaeddin Gürel, Benjamin J. Radford, Yaoyao Dai, Hansi Hettiarachchi, Niklas Stoehr, Tadashi Nomoto, Milena Slavcheva, Francielle Vargas, Aaqib Javid, Fatih Beyhan, and Erdem Yörük. 2022a. [Extended multilingual protest news detection - shared task 1, CASE 2021 and 2022](#). In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, pages 223–228, Abu Dhabi, United Arab Emirates (Hybrid). 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, Vanni Zavarella, Jakub Piskorski, Reyhan 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, Reyhan Yeniterzi, Osman Mutlu, and Erdem Yörük. 2022b. [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, Firat Duruşan, Çağrı Yoltar, Deniz Yuret, and Burak Gürel. 2021. [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).
- Guillaume Jacquet, Jakub Piskorski, Hristo Tanev, and Ralf Steinberger. 2019. Jrc tma-cc: Slavic named entity recognition and linking. participation in the bsnlp-2019 shared task. In *Proceedings of the 7th Workshop on Balto-Slavic Natural Language Processing*, pages 100–104.
- Avanthika K, Mrithula Kl, and Thenmozhi D. 2023. Ssn-nlp-ace@multimodal hate speech event detection 2023: Detection of hate speech and targets using logistic regression and svm. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Mohammad Kashif, Mohammad Zohair, and Saquib Ali. 2023. Lexical squad@multimodal hate speech event detection 2023: Multimodal hate speech detection

- using fused ensemble approach. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A robustly optimized bert pretraining approach. *arXiv:1907.11692*.
- Steve Fonin Mbouadeu, Martin Lorenzo, Ken Barker, and Oktie Hassanzadeh. 2023. An evaluation framework for mapping news headlines to event classes in a knowledge graph. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Osman Mutlu and Ali Hürriyetoğlu. 2023. Negative documents are positive: Improving event extraction performance using overlooked negative data. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Javier Osorio and Juan Vásquez. 2023. Classifying organized criminal violence in mexico using ml and llms. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Rajat Patel. 2023. Interosml @ causal news corpus 2023: Understanding causal relationships: Supervised contrastive learning for event classification. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Clionadh Raleigh, rew Linke, Håvard Hegre, and Joakim Karlsen. 2010. Introducing acled: An armed conflict location and event dataset. *Journal of peace research*, 47(5):651–660.
- 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 (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Timo Pierre Schrader, Simon Razniewski, Lukas Lange, and Annemarie Friedrich. 2023. Boschai @ causal news corpus 2023: Robust cause-effect span extraction using multi-layer sequence tagging and data augmentation. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Karanpreet Singh, Vajratiya Vajrobol, and Nitisha Aggarwal. 2023. IIC_Team@multimodal hate speech event detection 2023: Detection of hate speech and targets using xlm-roberta-base. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Milena Slavcheva, Hristo Tanev, and Onur Uca. 2023. On the road to a protest event ontology for bulgarian: Conceptual structures and representation design. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Fiona Anting Tan, Hansi Hettiarachchi, Ali Hürriyetoğlu, Tommaso Caselli, Onur Uca, Farhana Ferdousi Liza, and Nelleke Oostdijk. 2022a. [Event causality identification with causal news corpus - shared task 3, CASE 2022](#). In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, pages 195–208, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Fiona Anting Tan, Hansi Hettiarachchi, Ali Hürriyetoğlu, Nelleke Oostdijk, Tommaso Caselli, Tadashi Nomoto, Onur Uca, Farhana Ferdousi Liza, and See-Kiong Ng. 2023a. RECESS: Resource for extracting cause, effect, and signal spans. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, Bali, Indonesia. Association for Computational Linguistics.
- Fiona Anting Tan, Hansi Hettiarachchi, Ali Hürriyetoğlu, Nelleke Oostdijk, Onur Uca, Surendrabikram Thapa, and Farhana Ferdousi Liza. 2023b. [Event causality identification - shared task 3, case 2023](#). In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.
- Fiona Anting Tan, Ali Hürriyetoğlu, Tommaso Caselli, Nelleke Oostdijk, Tadashi Nomoto, Hansi Hettiarachchi, Iqra Ameer, Onur Uca, Farhana Ferdousi Liza, and Tiancheng Hu. 2022b. [The causal news corpus: Annotating causal relations in event sentences from news](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2298–2310, Marseille, France. European Language Resources Association.
- Hristo Tanev and Bertrand De Longueville. 2023. Where "where" matters : Event location disambiguation with a bert language model. In *Proceedings of*

the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE), Varna, Bulgaria (Hybrid). Association for Computational Linguistics.

Hristo Tanev, Jakub Piskorski, and Martin Atkinson. 2008. Real-time news event extraction for global crisis monitoring. In *Natural Language and Information Systems: 13th International Conference on Applications of Natural Language to Information Systems, NLDB 2008 London, UK, June 24-27, 2008 Proceedings 13*, pages 207–218. Springer.

Hristo Tanev, Nicolas Stefanovitch, Andrew Halterman, Onur Uca, Vanni Zavarella, Ali Hürriyetoğlu, Bertrand De Longueville, and Leonida Della Rocca. 2023a. Detecting and geocoding battle events from social media messages on the russo-ukrainian war: Shared task 2, case 2023. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*. Association for Computational Linguistics (ACL).

Hristo Tanev, Nicolas Stefanovitch, Andrew Halterman, Onur Uca, Vanni Zavarella, Ali Hürriyetoğlu, Bertrand De Longueville, and Leonida Della Rocca. 2023b. Where "where" matters : Event location identification with a BERT language model. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*. Association for Computational Linguistics (ACL).

Surendrabikram Thapa, Farhan Jafri, Ali Hürriyetoğlu, Francielle Vargas, Roy Ka-Wei Lee, and Usman Naseem. 2023. 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 (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.

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.

Elena Tuparova, Petar Ivanov, Andrey Tagarev, Svetla Boytcheva, and Ivan Koychev. 2023. Next: An event schema extension approach for closed-domain event extraction models. In *Proceedings of the 6th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE)*, Varna, Bulgaria (Hybrid). Association for Computational Linguistics.

Erdem Yörük, Ali Hürriyetoğlu, Fırat Duruşan, and Çağrı Yoltar. 2021. [Random sampling in corpus design: Cross-context generalizability in automated multicountry protest event collection](#). *American Behavioral Scientist*, 0(0):00027642211021630.