

# Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2022): Workshop and Shared Task Report

**Ali Hürriyetoglu**  
KNAW Humanities Cluster DHLab  
ali.hurriyetoglu@dh.huc.knaw.nl

**Hristo Tanev**  
European Commission  
hristo.tanev@ec.europa.eu

**Vanni Zavarella**  
University of Cagliari  
v.zavarella@unica.it

**Reyyan Yeniterzi**  
Sabanci University  
reyyan.yeniterzi@sabanciuniv.edu

**Osman Mutlu**  
Koc University  
omutlu@ku.edu.tr

**Erdem Yörük**  
Koc University  
eryoruk@ku.edu.tr

## Abstract

We provide a summary of the fifth edition of the CASE workshop that is held in the scope of EMNLP 2022. The workshop consists of regular papers, two keynotes, working papers of shared task participants, and task overview papers. This workshop has been bringing together all aspects of event information collection across technical and social science fields. In addition to the progress in depth, the submission and acceptance of multimodal approaches show the widening of this interdisciplinary research topic.

## 1 Introduction

The workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE) has become a significant venue where all technical and social science aspects of event information collection can be discussed in its fifth edition.<sup>1</sup> The 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP 2022) hosts the edition this year between December 7 and 11 in Abu Dhabi.<sup>2</sup>

Socio-political event extraction (SPE) has long been a challenge for the natural language processing (NLP) community as it requires sophisticated methods in defining event ontologies, creating language resources, and developing algorithmic approaches (Pustejovsky et al., 2003; Boroş, 2018; Chen et al., 2021). Social and political scientists have been working to create socio-political event (SPE) databases such as ACLED, EMBERS, GDELT, ICEWS, MMAD, PHOENIX, POLDEM, SPEED, TERRIER, and UCDP following similar steps for decades. These projects and the new ones increasingly rely on machine learning (ML), deep learning (DL), and NLP methods to deal better

<sup>1</sup><https://emw.ku.edu.tr/case-2022/>, accessed on November 14, 2022.

<sup>2</sup><https://2022.emnlp.org/>, accessed on November 14, 2022.

with the vast amount and variety of data in this domain (Hürriyetoglu et al., 2020; Hürriyetoglu et al., 2021b). Automation offers scholars not only the opportunity to improve existing practices but also to vastly expand the scope of data that can be collected and studied, thus potentially opening up new research frontiers within the field of SPEs, such as political violence and social movements. But automated approaches suffer from major issues like bias, generalizability, class imbalance, training data limitations, and ethical issues that have the potential to affect the results of automated text processing systems and their use drastically (Leins et al., 2020; Bhatia et al., 2020; Chang et al., 2019). Moreover, the results of the automated systems for SPE information collection have neither been comparable to each other nor been of sufficient quality (Wang et al., 2016; Schrodt, 2020).

Setting a clear path toward addressing these challenges is our main focus. We are confident that the program we put together for this year’s event after rigorous and thorough reviews, would bring us closer to that goal and beyond.

We provide a summary of the accepted papers in the following Section, which is Section 2. Next, the shared tasks that are organized in the scope of this workshop are described in Section 3. The keynote abstracts and invited talks are provided in sections 4 and 5. Finally, Section 6 conclude this report with a brief summary and future outlook.

## 2 Accepted papers

This year, out of 12 submissions 8 were accepted by the program committee. A quick summary of these papers are provided below.

- Thapa et al. (2022) releases a multimodal dataset that consists of of 5,680 text-image pairs of tweets and a baseline for hate speech detection in the context of Russia-Ukraine war.

- [You et al. \(2022\)](#) propose Event-Graph, a joint framework for event extraction, which encodes events as graphs. They represent event triggers and arguments as nodes in a semantic graph. Event extraction therefore becomes a graph parsing problem.
- [Desot et al. \(2022\)](#) suggests event and argument role detection as one task in a hybrid event detection approach and a novel method for automatic self-attention threshold selection.
- [Mehta et al. \(2022\)](#) cast socio-political conflict event extraction as a machine reading comprehension (MRC) task. In this approach, extraction of socio-political actors and targets from a sentence is framed as an extractive question answering problem conditioned on event type.
- [Raj et al. \(2022\)](#) propose a method that utilizes existing annotated unimodal data to perform event detection in another data modality using a zero-shot setting. They focus on protest detection in text and images, and show that a pretrained vision-and-language alignment model (CLIP) can be leveraged towards this end.
- [Sticha and Brenner \(2022\)](#) release a comprehensive, consolidated, and cohesive assassination dataset that is prepared utilizing a robust ML framework that prioritizes understandability through visualizations and generalizability through the ability to implement different ML algorithms.
- [Yaoyao et al. \(2022\)](#) train a model that can produce structured political event records at the sentence level. This approach is based on text-to-text sequence generation. They also describe a method for generating synthetic text and event record pairs that we use to fit a model.
- [Kiyet Akdemir and Hürriyetoğlu \(2022\)](#) approach the classification problem as an entailment problem and apply zero-shot ranking to socio-political texts. Documents that are ranked at the top can be considered positively classified documents and this reduces the close reading time for the information extraction process.

### 3 Shared tasks

Three tasks were organized in the scope of CASE 2022. Each one of these tasks shed light on a different aspect of event information collection. These are zero-shot and detailed multilingual event information, evaluation of state-of-the-art systems in replicating manually curated event datasets, and event causality detection.

#### 3.1 Task 1: Extended Multilingual Protest Event Detection

The extended multilingual protest news detection is the same shared task organized at CASE 2021 ([Hürriyetoğlu et al., 2021a](#)). This year we introduced additional data and languages at the evaluation stage.<sup>3</sup> This year, the Task 1 focused on evaluating the zero-shot prediction performance of the state-of-the-art systems for Subtask 1, document classification. The training set is the same with CASE 2021 data that is in English, Portuguese, and Spanish. But the evaluation data consists of the union of CASE 2021 test data and new data in both available and new languages. The new languages are Mandarin, Urdu, and Turkish. Details of CASE 2022 Task 1 is reported by [Hürriyetoğlu et al. \(2022\)](#).

#### 3.2 Task 2: Tracking COVID-19 protest events in the United States

This task aims at automatically replicating manually created event datasets. The participants of Task 1 are invited to run the systems they develop to tackle Task 1 on a news and a Twitter archive. This is a similar setting with the edition performed last year in the scope of CASE 2021 and reported by [Giorgi et al. \(2021\)](#). This year's results<sup>4</sup> are reported by [Zavarella et al. \(2022\)](#).

#### 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. This task focuses on the study of event causality in news, and therefore, introduces the Causal News Corpus ([Tan et al., 2022b](#)). The Causal News Corpus consists of 3,559 event sentences from CASE 2021 data, extracted from protest event

<sup>3</sup><https://github.com/emerging-welfare/case-2022-multilingual-event>, accessed on November 15, 2022.

<sup>4</sup>[https://github.com/zavavan/case2022\\_task2](https://github.com/zavavan/case2022_task2), accessed on November 14, 2022.

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. The two sub-tasks (Sequence Classification and Span Detection) work on the Causal News Corpus, and accurate, automated solutions are proposed for the detection and extraction of causal events in news. The detailed report of the task is provided in Tan et al. (2022a).<sup>5</sup>

## 4 Keynotes

Three scholars delivered two keynote speeches that are one on event extraction system development and one for error analysis of event information collection systems. The speakers were invited according to our tradition of having one keynote with technical and another one with social and political sciences, background. We provide abstracts of the keynote speeches as they are provided by the keynote speakers in the following subsections. Section 4.1 and Section 4.2 are contributions of Prof. Thien Huu Nguyen<sup>6</sup> and Prof. Scott Althaus<sup>7</sup> and Prof. J. Craig Jenkins<sup>8</sup> respectively.<sup>9</sup>

### 4.1 Event Extraction in the Era of Large Language Models: Structure Induction and Multilingual Learning

Events such as protests, disease outbreaks, and natural disasters are prevalent in text from different languages and domains. Event Extraction (EE) is an important task of Information Extraction that aims to identify events and their structures in unstructured text. The last decade has witnessed significant progress for EE research, featuring deep learning and large language models as the state-of-the-art technologies. However, a key issue of existing EE methods involves modeling input text sequentially to solve each EE tasks separately, thus limiting the abilities to encode long text and capture various types of dependencies to improve EE performance. In this talk, I will present some of

<sup>5</sup><https://github.com/tanfiona/CausalNewsCorpus>, accessed on November 14, 2022.

<sup>6</sup><https://ix.cs.uoregon.edu/~thien/>, accessed on November 14, 2022.

<sup>7</sup><https://ix.cs.uoregon.edu/~thien/>, accessed on November 14, 2022.

<sup>8</sup><https://sociology.osu.edu/people/jenkins.12>, accessed on November 15, 2022.

<sup>9</sup>The personal pronoun usage such as ‘I’ and ‘we’ in the following subsections indicate the keynote speakers and not the authors of this report.

our recent efforts to address this issue where text structures are explicitly learned to realize important objects and their interactions to facilitate the predictions for EE.

In addition, current EE research still mainly focuses on a few popular languages, e.g., English, Chinese, Arabic, and Spanish, leaving many other languages unexplored for EE. In this talk, I will also introduce our current research focus on developing evaluation benchmarks and models to extend EE systems to multiple new languages, i.e., multilingual and cross-lingual learning for EE. Finally, I will highlight some research challenges that can be studied in future work for EE.

### 4.2 A total error approach to validating event data that is transparent, scalable, and practical to implement

There are at least two reasonable ways to make your way toward where you want to go: looking down to carefully place one foot in front of the other, and looking up to focus on where you hope to arrive. Looking up beats looking down if there’s a particular destination in mind, and for constructing valid event data that destination usually takes the form of high-quality human judgment. Yet many approaches to generating event data on protests and acts of political violence using fully-automated systems implicitly adopt a “looking down” approach by benchmarking validity as a series of incremental improvements over prior algorithmic efforts. And even those efforts that adopt a “looking up” approach often treat human-generated gold standard data as if it was prima facie valid, without ever testing or confirming the accuracy of this assumption. It stands to reason that if we want to automatically produce valid event data that approaches the validity of human judgment, then we also need to validate the human judgment tasks that provide the point of comparison. But because of obvious difficulties in implementing such a rigorous assessment within the time and budget constraints of typical research projects, this more rigorous double-validation approach is rarely attempted.

This presentation outlines a “looking up” approach for double-validating fully-automated event data developed by the Cline Center for Advanced Social Research at the University of Illinois Urbana-Champaign (USA), illustrates that approach with a test of the precision and recall for two widely-used event classification systems (the

PETRARCH-2 coder used in Phoenix and TERRIER, as well as the BBN ACCENT coder used in W-ICEWS), and demonstrates the utility of the approach for developing fully-automated event data algorithms with levels of validity that approach the quality of human judgment.

The first part of the talk reviews the Cline Center’s total error framework for identifying 19 types of error that can affect the validity of event data and addresses the challenge of applying a total error framework when authoritative ground truth about the actual distribution of relevant events is lacking (Althaus et al., 2022). We argue that carefully constructed gold standard datasets can effectively benchmark validity problems even in the absence of ground truth data about event populations. We propose that a strong validity assessment for event data should, at a minimum, possess three characteristics. First, there should be a standard describing ideal data; a gold standard that, in the best case, takes the form of ground truth. Second, there should be a direct “apples to apples” comparison of outputs from competing methods given identical input. Third, the test should use appropriate metrics for measuring agreement between the gold standard and data produced by competing approaches.

The second part of the talk presents the results of a validation exercise meeting all three criteria that is applied to two algorithmic event data pipelines: the Python Engine for Text Resolution and Related Coding Hierarchy (PETRARCH-2) and the BBN ACCENT event coder. It then reviews a recent Cline Center project that has built a fully-automated event coder which produces dramatic improvements in validity over both PETRARCH-2 and BBN ACCENT by leveraging the total error framework and a reliance on the double-validation approach using high-quality gold standard benchmark datasets.

## 5 Invited talks

Papers that are accepted to be published in the Findings of EMNLP 2022 and related to our workshop theme were invited to be presented during our workshop. The authors of the following papers were invited for presenting their papers:

- Jiao et al. (2022) define the task open-vocabulary argument role prediction. The goal of this task is to infer a set of argument roles for a given event type. They propose a novel unsupervised framework, ROLEPRED for this

task and release a new human-annotated event extraction dataset including 139 customized argument roles with rich semantics.

- Faghihi et al. (2022) presents CrisisLTL-Sum, the largest dataset of local crisis event timelines about wildfires, local fires, traffic, and storms available to date. CrisisLTLSum was built using a semi-automated cluster-then-refine approach to collect data from the public Twitter stream.
- Ding et al. (2022) design a neural model that they refer to as the Explicit Role Interaction Network (ERIN) which allows for dynamically capturing the correlations between different argument roles within an event.
- Gao et al. (2022) present Mask-then-Fill, a flexible and effective data augmentation framework for event extraction. This approach allows for more flexible manipulation of text and thus can generate more diverse data while keeping the original event structure unchanged.

## 6 Conclusion

The CASE workshop series has been contributing to both technical advancement in terms of shared task organization and being a venue for scholars working at the intersection of social sciences and event extraction. This role become more significant as these series are known to a wider community. Following steps of this series should serve the community by preserving its interdisciplinary setting, welcoming new methodologies, and promoting responsible development and utilization of the results of this scholarship.

## References

- Scott Althaus, Buddy Peyton, and Dan Shalmon. 2022. [A total error approach for validating event data](#). *American Behavioral Scientist*, 66(5):603–624.
- Shraey Bhatia, Jey Han Lau, and Timothy Baldwin. 2020. You are right. i am alarmed—but by climate change counter movement. *arXiv preprint arXiv:2004.14907*.
- Emanuela Boros. 2018. *Neural Methods for Event Extraction*. Ph.D. thesis, Université Paris-Saclay.
- Kai-Wei Chang, Vinod Prabhakaran, and Vicente Ordonez. 2019. [Bias and fairness in natural language processing](#). In *Proceedings of the 2019 Conference*

- on *Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP): Tutorial Abstracts*, Hong Kong, China. Association for Computational Linguistics.
- Muhao Chen, Hongming Zhang, Qiang Ning, Manling Li, Heng Ji, Kathleen McKeown, and Dan Roth. 2021. Event-centric natural language understanding.
- Thierry Desot, Orphee De Clercq, and Veronique Hoste. 2022. A hybrid knowledge and transformer-based model for event detection with automatic self-attention threshold, layer and head selection. In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2022)*, online. Association for Computational Linguistics (ACL).
- Nan Ding, Chunming Hu, Kai Sun, Samuel Mensah, and Richong Zhang. 2022. Explicit role interaction network for event argument extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, Online. Association for Computational Linguistics.
- Hossein Rajaby Faghihi, Bashar Alhafni, Ke Zhang, Shihao Ran, Joel Tetreault, and Alejandro Jaimes. 2022. CrisisItsum: A benchmark for local crisis event timeline extraction and summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, Online. Association for Computational Linguistics.
- Jun Gao, Changlong Yu, Wei Wang, and Ruifeng Zhao, Huan Xu. 2022. Mask-then-fill: A flexible and effective data augmentation framework for event extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, Online. Association for Computational Linguistics.
- Salvatore Giorgi, Vanni Zavarella, Hristo Tanev, Nicolas Stefanovitch, Sy Hwang, Hansi Hettiarachchi, Tharindu Ranasinghe, Vivek Kalyan, Paul Tan, Shaun Tan, Martin Andrews, Tiancheng Hu, Niklas Stoehr, Francesco Ignazio Re, Daniel Vegh, Dennis Atzenhofer, Brenda Curtis, and Ali Hürriyetoğlu. 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, Online. Association for Computational Linguistics.
- Ali Hürriyetoğlu, Osman Mutlu, Firat Duruşan, Onur Uca, Alaeddin Selçuk Gürel, Benjamin Radford, Yaoyao Dai, Hansi Hettiarachchi, Niklas Stoehr, Tadashi Nomoto, Milena Slavcheva, Francielle Vargas, Aaqib Javid, Fatih Beyhan, and Erdem Yörük. 2022. 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 2022)*, online. Association for Computational Linguistics (ACL).
- 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, 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).
- Yizhu Jiao, Sha Li, Yiqing Xie, Ming Zhong, Heng Ji, and Jiawei Han. 2022. Open-vocabulary argument role prediction for event extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, Online. Association for Computational Linguistics.
- Kıymet Kıymet Akdemir and Ali Hürriyetoğlu. 2022. Zero-shot ranking socio-political texts with transformer language models to reduce close reading time. In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2022)*, online. Association for Computational Linguistics (ACL).
- Kobi Leins, Jey Han Lau, and Timothy Baldwin. 2020. [Give me convenience and give her death: Who should decide what uses of NLP are appropriate, and on what basis?](#) In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2908–2913, Online. Association for Computational Linguistics.
- Sneha Mehta, Huzefa Rangwala, and Naren Ramakrishnan. 2022. Improving zero-shot event extraction via sentence simplification. In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2022)*, online. Association for Computational Linguistics (ACL).
- James Pustejovsky, José M Castano, Robert Ingria, Roser Sauri, Robert J Gaizauskas, Andrea Setzer, Graham Katz, and Dragomir R Radev. 2003. Timeml: Robust specification of event and temporal expressions in text. *New directions in question answering*, 3:28–34.

- Ria Raj, Kajsa Andreasson, and Tobias Norlund. 2022. Cross-modal transfer between vision and language for protest detection. In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2022)*, online. Association for Computational Linguistics (ACL).
- Philip A. Schrod. 2020. [Keynote abstract: Current open questions for operational event data](#). In *Proceedings of the Workshop on Automated Extraction of Socio-political Events from News 2020*, page 8, Marseille, France. European Language Resources Association (ELRA).
- Abigail Sticha and Paul R. Brenner. 2022. Hybrid knowledge engineering leveraging a robust ml framework to produce an assassination dataset. In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2022)*, online. Association for Computational Linguistics (ACL).
- Fiona Anting Tan, Ali Hürriyetoğlu, Tommaso Caselli, Nelleke Oostdijk, Hansi Hettiarachchi, Tadashi Nomoto, Onur Uca, and Farhana Ferdousi Liza. 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 2022)*, Online. 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 Language Resources and Evaluation Conference*, pages 2298–2310, Marseille, France. European Language Resources Association.
- 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 *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2022)*, online. Association for Computational Linguistics (ACL).
- Wei Wang, Ryan Kennedy, David Lazer, and Naren Ramakrishnan. 2016. [Growing pains for global monitoring of societal events](#). *Science*, 353(6307):1502–1503.
- Dai Yaoyao, Benjamin Radford, and Andrew Halterman. 2022. Political event coding as text-to-text sequence generation. In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2022)*, online. Association for Computational Linguistics (ACL).
- Huiling You, David Samuel, Samia Touileb, and Lilja Øvrelid. 2022. Eventgraph: Event extraction as semantic graph parsing. In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2022)*, online. Association for Computational Linguistics (ACL).
- Vanni Zavarella, Hristo Tanev, Ali Hürriyetoğlu, Peratham Wiriathamabhum, and Bertrand De Longueville. 2022. Tracking COVID-19 protest events in the United States. Shared Task 2: Event Database Replication, CASE 2022. In *Proceedings of the 5th Workshop on Challenges and Applications of Automated Extraction of Socio-political Events from Text (CASE 2022)*, online. Association for Computational Linguistics (ACL).