

Automated Extraction of Socio-political Events from News (AESPEN): Workshop and Shared Task Report

Ali Hürriyetoglu*, Vanni Zavarella†, Hristo Tanev†, Erdem Yörük*, Ali Safaya*, Osman Mutlu*

*Koç University,
Rumelifeneri Yolu 34450, Sarıyer, İstanbul/Turkey
{ahurriyetoglu, eryoruk, asafaya19, omutlu}@ku.edu.tr

†European Commission Joint Research Centre
Via E. Fermi, 2749, 21027 Ispra VA, Italy
{vanni.zavarella, hristo.tanev}@ec.europa.eu

Abstract

We describe our effort on automated extraction of socio-political events from news in the scope of a workshop and a shared task we organized at Language Resources and Evaluation Conference (LREC 2020). We believe the event extraction studies in computational linguistics and social and political sciences should further support each other in order to enable large scale socio-political event information collection across sources, countries, and languages. The event consists of regular research papers and a shared task, which is about event sentence coreference identification (ESCI), tracks. All submissions were reviewed by five members of the program committee. The workshop attracted research papers related to evaluation of machine learning methodologies, language resources, material conflict forecasting, and a shared task participation report in the scope of socio-political event information collection. It has shown us the volume and variety of both the data sources and event information collection approaches related to socio-political events and the need to fill the gap between automated text processing techniques and requirements of social and political sciences.

Keywords: socio-political events, information extraction, event extraction, machine learning, natural language processing, computational linguistics, social sciences, political sciences

1. Introduction

Automatic construction of socio-political event databases has long been a challenge for the natural language processing (NLP) and social and political science communities in terms of algorithmic approaches and language resources required to develop automated tools (Chenoweth and Lewis, 2013; Weidmann and Rød, 2019; Raleigh et al., 2010). At the same time, social and political scientists have been working on creating socio-political event databases for decades using manual (Yoruk, 2012), semi-automatic (Nardulli et al., 2015), and automatic approaches (Leetaru and Schrodt, 2013; Boschee et al., 2013; Schrodt et al., 2014; Sönmez et al., 2016). However, the results yielded by these approaches to date are either not of sufficient quality or require tremendous effort to be replicated on new data (Wang et al., 2016; Ward et al., 2013; Etinger et al., 2017). On the one hand, manual or semi-automatic methods require high-quality human effort; on the other hand, state-of-the-art automated event detection systems are not accurate enough for their output to be used directly without human moderation.

The NLP community has achieved some consensus on the treatment of events both in terms of task definition and appropriate techniques for their detection (Pustejovsky et al., 2005; Doddington et al., 2004; Song et al., 2015; Getman et al., 2018). However, in order to be useful, these formalisms and related systems need to be adjusted or extended for each type of event in relation to certain use cases. The social and political scientists spend a similar effort for formalising event types such as (CAMEO) (Gerner et al., 2002) and implement the aforementioned systems that vary

from rule-based to fully automatic approaches. Unfortunately, any new project in this line still finds itself making design decisions such as using only the heading sentences in a news article or not considering coreference information (Boschee et al., 2013) without being able to quantify their effect. Therefore, we think these communities should investigate ways of supporting each other in order to reach a consensus and enable any prospective event information collection project as robustly and predictably as possible. Given the aforementioned limitations, there is an increasing tendency to rely on machine learning (ML) and NLP methods to deal better with the vast amount and variety of data to be processed. Consequently, we thought it was time to held a workshop on Automated Extraction of Socio-political Events from News (AESPEN)¹ at Language Resources and Evaluation Conference (LREC 2020).² The purpose of this workshop was to inspire the emergence of innovative technological and scientific solutions in the field of event detection and event metadata extraction from news, as well as the development of evaluation metrics for socio-political event recognition. Moreover, the workshop aimed at triggering a deeper understanding of the usability of socio-political event datasets.

We organized a shared task as a continuation of the Conference and Labs of the Evaluation Forum (CLEF 2019) task ProtestNews (Hürriyetoglu et al., 2019a; Hürriyetoglu et al., 2019b), which was on cross-context document clas-

¹<https://emw.ku.edu.tr/aespen-2020/>, accessed on April 18, 2020.

²<https://lrec2020.lrec-conf.org/>, accessed on April 18, 2020.

sification, event sentence detection, and event extraction pertaining to protest events. We aimed at establishing a benchmark for the event sentence coreference identification (ESCI) sub-task within the scope of the AESPEN workshop. The scope of this shared task was on clustering given event related sentences so that each cluster consists of sentences about the same event.

We provide details of our motivation in Section 2. Then, we introduce the ESCI shared task in Section 3. Finally, we briefly describe the accepted papers and the shared task results in Section 4. We conclude this report in section 5.

2. Motivation

Automating political event collection requires the availability of gold-standard corpora that can be used for system development and evaluation. Moreover, automated tool performances need to be reproducible and comparable. Although a tremendous effort is being spent on creating socio-political event databases such as ACLED (Raleigh et al., 2010), the Global Database of Events, Language, and Tone (GDELT) (Leetaru and Schrodt, 2013), the Mass Mobilization on Autocracies Database (MMAD) (Weidmann and Rød, 2019), the Integrated Crisis Early Warning System (ICEWS) (Boschee et al., 2013), and the Protest Dataset 30 European countries (PolDem) (Kriesi et al., 2019) we believe there is still a lot of room for improvement and harmonisation of the event schemas and tasks. This limitation causes the definition of the events and automated event information collection tool performances to be restricted to single projects. Consequently, the lack of comparable and reproducible settings hinders progress on this task.

We invited contributions from researchers in NLP, ML and Artificial Intelligence (AI) involved in automated event data collection, as well as researchers in social and political sciences, conflict analysis and peace studies, who make use of this kind of data for their analytical work. Our goal was to enable the emergence of innovative NLP and information extraction (IE) solutions that can deal with the current stream of information, manage the risks of information overload, identify different sources and perspectives, and provide unitary and intelligible representations of the larger and long-term storylines behind news articles.

Our workshop provided a venue for discussing the creation and facilitation of language resources in the social and political sciences domain. Social and political scientists were interested in reporting and discussing the automated tools and comparing traditional coding approaches with automated tools. Computational linguistics and machine learning practitioners and researchers benefited from being challenged by real-world use cases, in terms of event data extraction, representation and aggregation.

We invited work on all aspects of automated coding of socio-political events from monolingual or multilingual news sources. This includes (but is not limited to) the following topics: event metadata extraction, source bias mitigation, event data schema and representation, event information duplication detection, extracting events beyond a sentence in a document, training data collection and annotation processes, event coreference (in- and cross-document), sub-event and event subset relations, event dataset evalu-

ation and validity metrics, event datasets quality assessments, defining, populating and facilitating event ontologies, automated tools for relevant subtasks, understanding the limits that are introduced by copyright rules and ethical concerns and ethical design.

3. Shared Task

A news article may contain one or more events that are expressed with one or more sentences. Identifying event sentences that are about the same event is necessary in order to collect event information robustly. Therefore, we should develop methods that are able to identify whether a group of sentences are about the same event. Reliable identification of this relation will enable us to determine how many events are reported in a news article as well. Moreover, solving this problem has the potential to facilitate cross-document event sentence relation identification in the long term. Therefore, we should develop methods that are able to identify whether a group of event sentences are about the same event. Consequently, we organized the ESCI shared task in the hopes of attracting attention to this problem and possibly provide a benchmark for it.

We examined our gold standard corpus that contains 1,290 events in 712 documents annotated at token level for their event information (Hürriyetoğlu et al., 2019a; Hürriyetoğlu et al., 2019b; Hürriyetoğlu et al., 2020). These documents are the positively labelled instances of random samples and active learning based samples based on these random samples. We have observed that 60% of the news articles contain information about a single event. The remaining documents contain information about multiple events, which sums up to 45% of the total event count. Only 45% of the events are expressed with only a single sentence.

Consequently, we think protest event collection systems should take these phenomena into account and introduce the ESCI shared task. As training data participants of the data challenge received event related sentences and their true clustering in a news article, in which a cluster represents all sentences about an event. This data was extracted from 404 documents. The documents that contain a single event sentence were excluded from this exercise, since there is only one possible clustering in that case. The number of events per document in the training data is 1 for 207 and 2 for 132 documents. The remaining 65 documents contain 3 or more events. The task of the participants was to develop systems that can predict grouping of the given sentences that consists of events on test data, extracted from 100 documents, and that was delivered to them one week before the deadline. The correct grouping of the test set was not shared with the participants. The evaluation metric is Adjusted Rand Index (ARI) as implemented by Scikit-learn (Hubert and Arabie, 1985).³ We calculated macro and micro versions of this score. The macro version calculates average of the per document scores from all of the documents independent of how many event sentences are there in each document. However, the micro score weights the

³https://scikit-learn.org/stable/modules/generated/sklearn.metrics.adjusted_rand_score.html, accessed on April 19, 2020.

per document score with the number of the event sentences in a document. We report the F1 score that is calculated similarly as well.

The event type is a protest in the scope of this task. The event we simply refer to as protest events are comprised within the scope of contentious politics. Contentious politics events refers to politically motivated collective action events which lay outside the official mechanisms of political participation associated with formal government institutions of the country in which the said action takes place.⁴ The data is shared with the researchers, who signed an application form that limits the use of the data only for research purposes, as a file that contains lines of JSON objects. Each JSON object contain all event sentences that are identified in a news article and their clustering, which is found in the *event_clusters* field.

A sample JSON object is presented below. The *url* field is provided only as an ID. The numbers in the *sentence_no* correspond to sentences in the *sentences* field in the same order. The *event_clusters* field provides the correct clustering of the event sentences. For instance, below in Listing 1, the first and third sentences are about the same event. But, the second sentence is about a separate event.

Listing 1: Event sentences that are extracted from a document in the order they occur in a sentence.

```

1 {
2   "url": "http://www.newindianexpress.
   com/nation/2009/aug/25/congress-
   demands-advanis-apology-80257",
3   "sentences": [
4     "Singh had recently blamed Advani
   for coming to Gujarat Chief
   Minister Narendra Modi ' s
   rescue and ensured that he was
   not sacked , in the wake of the
   riots .",
5     "On Kandahar plane hijack issue ,
   Singh said Advani was not
   speaking the truth .",
6     "Elaborating on the three issues ,
   Singhvi said , The BJP gave
   sermons on Raj Dharma and
   turned a Nelson ' s eye to the
   communal carnage , which became
   a big blot on the fair name of
   the country ."
7   ],
8   "sentence_no": [4, 6, 14],
9   "event_clusters": [[4, 14], [6]]
10 }

```

We have calculated three baseline scores on the test data. First, we checked score of a dummy predictor that assigns all event sentences to a single cluster all the time, i.e., min-

⁴You can find detailed information about how a protest is defined and how event sentences are labelled on our annotation manual, which is on https://github.com/emerging-welfare/general_info/tree/master/annotation-manuals.

imum cluster prediction (MinC). Second, another dummy baseline predicts as each event sentence as being in a separate cluster in a document, i.e., maximum cluster prediction (MaxC). Finally, we used BERT sentence representations (Devlin et al., 2019) to train a multilayer perceptron (MLP) model that i) first evaluates each possible sentence pair in the document, ii) then assign a positive or negative label indicating that this pair of sentences is co-referent, iii) finally using the correlation clustering algorithm (Bansal et al., 2004) we take those labeled pairs and cluster them.⁵ The scores of these methods are provided in Table 1 as *MinC*, *MaxC*, and *MLP*. The slightly low scores obtained from the dummy systems direct us to use the MLP system as the baseline we share with the participants. Note that the strength of the dummy baselines changes according to data distribution in test data.

	ARI		F1	
	Macro	Micro	Macro	Micro
MinC	.5000	.4040	.5000	.4040
MaxC	.1071	.0628	.3476	.3722
MLP	.5077	.4064	.5560	.4840

Table 1: Adjusted Random Index (ARI) and F1 for each baseline system.

4. Submissions

The workshop has attracted nine papers as regular paper submissions and one as a shared task participation report. The shared task report and seven of the regular papers were accepted on the basis of the reviews, which were five per paper, performed by the program committee members.

The accepted regular papers can be grouped as i) evaluation of state-of-the-art machine learning approaches by Büyüköz et al. (2020), Olsson et al. (2020), and Piskorski and Jacquet (2020), ii) introduction of a new data set by Radford (2020), iii) projects of event information collection by Osorio et al. (2020) and Papanikolaou and Papanikolaou (2020), and iv) forecasting of political conflict by Halkia et al. (2020).

The evaluation of Büyüköz et al. (2020) and Olsson et al. (2020) show that state-of-the-art deep learning models such as BERT and ELMo (Peters et al., 2018) yield consistently higher performance than traditional ML methods such as support vector machines (SVM) on conflict and protest event data respectively. Piskorski and Jacquet (2020) have found that TF-IDF weighted character n-gram based SVM model performs better than an SVM model that uses pre-trained embeddings such as GLOVE (Pennington et al., 2014), BERT, and FASTTEXT (Mikolov et al., 2018) in most of the experiments on conflict data.

Radford (2020) introduces the dataset *Headlines of War* for cross-document coreference resolution for the news headlines. The dataset consists of positive samples from *Militarized Interstate Disputes* dataset and negative samples

⁵The code for this system is available on <https://github.com/alisafrica/event-coreference>, accessed on April 21, 2020.

from New York Times.⁶ The description of this invaluable resource is accompanied with a detailed discussion of its utility and caveats.

Osorio et al. (2020) introduce Hadath that is a supervised protocol for event information collection from Arabic sources. The utility of Hadath was demonstrated in processing news reported between 2012 and 2012 in Afghanistan. In the scope of the other event information collection study, Papanikolaou and Papageorgiou1 (2020) processed two news sources in Greek from Greece to create a database of protest events for the period between 1996 and 2014. Osorio et al and Papanikolaou and Papageorgiou utilized fully automatic tools that integrate supervised machine learning and rule based methodologies at various degrees.

Finally, Halkia et al. (2020) presents a material conflict forecasting study that exploits available event databases GDELT and ICEWS. Their results demonstrate that it is possible to correctly predict social upheaval using the methodology they propose, which utilizes Long-Short Term Memory (LSTM) (Hochreiter and Schmidhuber, 1997).

We have received expression of interest from 12 research teams, of which 6 teams signed the application form and received the data. Two of these teams sent their predictions on the test data. The scores of these methods are illustrated in Table 2. Finally, only Örs et al. (2020) submitted a paper about their work. This team reported their work as consisting of three steps. First, they use a transformer based model, which is ALBERT (Lan et al., 2020), to predict whether a pair of sentences refer to the same event or not. Later, they use these predictions as the initial scores and recalculate the pair scores by considering the relation of sentences in a pair with respect to other sentences. As the last step, final scores between these sentences are used to construct the clusters, starting with the pairs with the highest scores.

	ARI		F1	
	Macro	Micro	Macro	Micro
Örs et al.	.6006	.4644	.6736	.5898
UNC Charlotte	.3388	.3253	.4352	.3284

Table 2: Adjusted Random Index (ARI) and F1 for each baseline system.

5. Concluding Remarks

We have provided a brief summary of the workshop Automated Extraction of Socio-political Events from News (AESPEN) and the shared task Event Sentence Coreference Identification (ESCI) we organized in the scope of Language Resources and Evaluation Conference (LREC 2020). The variety of the submitted papers show that we could bring the ML, NLP, and social and political science communities together. Although the breadth of the topics were limited, the technical depth and timeliness of the contributions show that the workshop contribute to the discipline

⁶<https://spiderbites.nytimes.com>, accessed on April 21, 2020.

of automatic extraction of socio-political events. The papers about processing Arabic and Greek sources are significant contributions to the understanding of how should we handle languages other than English. Finally, the shared task ESCI demonstrated the prevalence of the event coreferences, some baselines for handling them, and a state-of-the-art system that is able to tackle this task.

We consider this workshop as a beginning. We expect this effort to be extended both in depth and in breadth since we think the work presented is only the tip of the iceberg considering the recent projects and technical potential introduced by deep learning technologies.

6. Acknowledgements

The authors from Koç University are funded by the European Research Council (ERC) Starting Grant 714868 awarded to Dr. Erdem Yörük for his project Emerging Welfare. We appreciate contributions of the program committee members, who are in alphabetical order Svetla Boycheva, Firat Duruşan, Theresa Gessler, Christian Göbel, Burak Gürel, Matina Halkia, Sophia Hunger, J. Craig Jenkins, Liron Lavi, Jasmine Lorenzini, Bernardo Magnini, Osman Mutlu, Nelleke Oostdijk, Arzucan Özgür, Jakub Piskorski, Lidia Pivovarova, Benjamin J. Radford, Clionadh Raleigh, Ali Safaya, Parang Saraf, Philip Schrod, Manuela Speranza, Aline Villavicencio, Çağrı Yoltar, Kalliopi Zervanou and of the keynote speaker Clionadh Raleigh. We are grateful to the management of the Competence Centre on Text Mining and Analysis (CC-TMA) at European Commission Joint Research Center (JRC) for the support. Any opinions, findings, conclusions, or suggestions expressed here are those of the authors and do not necessarily reflect the view of the sponsor(s) or authors' employer(s).

7. Bibliographical References

- Bansal, N., Blum, A., and Chawla, S. (2004). Correlation clustering. *Mach. Learn.*, 56(1-3):89–113, June.
- Boschee, E., Natarajan, P., and Weischedel, R. (2013). Automatic Extraction of Events from Open Source Text for Predictive Forecasting. In V.S. Subrahmanian, editor, *Handbook of Computational Approaches to Counterterrorism*, pages 51–67. Springer New York, New York, NY.
- Büyükoğuz, B., Hürriyetoğlu, A., and Özgür, A. (2020). Analyzing ELMo and DistilBERT on Socio-political News Classification. In *Proceedings of the Workshop Automated Extraction of Socio-political Events from News (AESPEN)*.
- Chenoweth, E. and Lewis, O. A. (2013). Unpacking non-violent campaigns: Introducing the NAVCO 2.0 dataset. *Journal of Peace Research*, 50(3):415–423.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (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, June. Association for Computational Linguistics.

- Doddington, G., Mitchell, A., Przybocki, M., Ramshaw, L., Strassel, S., and Weischedel, R. (2004). The automatic content extraction (ACE) program – tasks, data, and evaluation. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal, May. European Language Resources Association (ELRA).
- Ettlinger, A., Rao, S., Daumé III, H., and Bender, E. M. (2017). Towards Linguistically Generalizable NLP Systems: A Workshop and Shared Task. In *Proceedings of the First Workshop on Building Linguistically Generalizable NLP Systems*, pages 1–10. Association for Computational Linguistics.
- Gerner, D. J., Schrodt, P. A., Yilmaz, O., and Abu-Jabr, R. (2002). Conflict and mediation event observations (cameo): A new event data framework for the analysis of foreign policy interactions. *International Studies Association, New Orleans*.
- Getman, J., Ellis, J., Strassel, S., Song, Z., and Tracey, J. (2018). Laying the groundwork for knowledge base population: Nine years of linguistic resources for TAC KBP. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May. European Language Resources Association (ELRA).
- Halkia, M., Ferri, S., Papazoglou, M., van Damme, M.-S., and Thomakos, D. (2020). Conflict Event Modelling: Research Experiment and Event Data Limitations. In *Proceedings of the Workshop Automated Extraction of Socio-political Events from News (AESPEN)*.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8):1735–1780.
- Hubert, L. and Arabie, P. (1985). Comparing partitions. *Journal of classification*, 2(1):193–218.
- Hürriyetoğlu, A., Yörük, E., Yüret, D., Yoltar, Ç., Gürel, B., Duruşan, F., and Mutlu, O. (2019a). A task set proposal for automatic protest information collection across multiple countries. In Leif Azzopardi, et al., editors, *Advances in Information Retrieval*, pages 316–323, Cham. Springer International Publishing.
- Hürriyetoğlu, A., Yörük, E., Yüret, D., Yoltar, Ç., Gürel, B., Duruşan, F., Mutlu, O., and Akdemir, A. (2019b). Overview of clef 2019 lab protestnews: Extracting protests from news in a cross-context setting. In Fabio Crestani, et al., editors, *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, pages 425–432, Cham. Springer International Publishing.
- Hürriyetoğlu, A., Yörük, E., Yüret, D., Mutlu, O., Yoltar, Ç., Gürel, B., and Duruşan, F. (2020). Cross-context news corpus for protest events related knowledge base construction. In *Automated Knowledge Base Construction (AKBC)*, June.
- Kriesi, H., Wüest, B., Lorenzini, J., Makarov, P., Enggist, M., Rothenhäusler, K., Kurer, T., Häusermann, Silja, P. W., Altıparmakı, A., Borbáth, E., Bremer, B., Gessler, T., Hunger, S., Hutter, S., Schulte-Cloos, J., and Wang, C. (2019). PolDem – Protest Event Dataset 30.
- Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., and Soricut, R. (2020). Albert: A lite bert for self-supervised learning of language representations. In *International Conference on Learning Representations*.
- Leetaru, K. and Schrodt, P. A. (2013). GDELT: Global data on events, location, and tone, 1979–2012. In *ISA annual convention*, volume 2, pages 1–49. Citeseer.
- Mikolov, T., Grave, E., Bojanowski, P., Puhersch, C., and Joulin, A. (2018). Advances in pre-training distributed word representations. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May. European Language Resources Association (ELRA).
- Nardulli, P. F., Althaus, S. L., and Hayes, M. (2015). A Progressive Supervised-learning Approach to Generating Rich Civil Strife Data. *Sociological Methodology*, 45(1):148–183.
- Olsson, F., Sahlgren, M., Abdesslem, F. B., Ekgren, A., and Eck, K. (2020). Text Categorization for Conflict Event Annotation. In *Proceedings of the Workshop Automated Extraction of Socio-political Events from News (AESPEN)*.
- Örs, K. F., Yeniterzi, S., and Yeniterzi, R. (2020). Event Clustering within News Articles. In *Proceedings of the Workshop Automated Extraction of Socio-political Events from News (AESPEN)*.
- Osorio, J., Reyes, A., Beltran, A., and Ahmadzai, A. (2020). Supervised Event Coding from Text Written in Arabic: Introducing Hadath. In *Proceedings of the Workshop Automated Extraction of Socio-political Events from News (AESPEN)*.
- Papanikolaou, K. and Papageorgiou, H. (2020). Protest Event Analysis: A Longitudinal Analysis for Greece. In *Proceedings of the Workshop Automated Extraction of Socio-political Events from News (AESPEN)*.
- Pennington, J., Socher, R., and Manning, C. (2014). Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar, October. Association for Computational Linguistics.
- Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. (2018). Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Piskorski, J. and Jacquet, G. (2020). TF-IDF Character Ngrams versus Word Embedding-based Models for Fine-grained Event Classification: A Preliminary study. In *Proceedings of the Workshop Automated Extraction of Socio-political Events from News (AESPEN)*.
- Pustejovsky, J., Knippen, R., Littman, J., and Saurí, R. (2005). Temporal and event information in natural language text. *Language resources and evaluation*, 39(2-3):123–164.
- Radford, B. (2020). Seeing the Forest and the Trees: Detection and Cross-Document Coreference Resolution of Militarized Interstate Disputes. In *Proceedings of*

- the Workshop Automated Extraction of Socio-political Events from News (AESPEN).*
- Raleigh, C., Linke, A., Hegre, H., and Karlsen, J. (2010). Introducing acled: an armed conflict location and event dataset: special data feature. *Journal of peace research*, 47(5):651–660.
- Schrodt, P. A., Beielser, J., and Idris, M. (2014). Three’s a charm?: Open event data coding with el: Diablo, Petrarch, and the open event data alliance. In *ISA Annual Convention*.
- Song, Z., Bies, A., Strassel, S., Riese, T., Mott, J., Ellis, J., Wright, J., Kulick, S., Ryant, N., and Ma, X. (2015). From light to rich ERE: Annotation of entities, relations, and events. In *Proceedings of the The 3rd Workshop on EVENTS: Definition, Detection, Coreference, and Representation*, pages 89–98, Denver, Colorado, June. Association for Computational Linguistics.
- Sönmez, Ç., Özgür, A., and Yörük, E. (2016). Towards building a political protest database to explain changes in the welfare state. In *Proceedings of the 10th SIGHUM Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, pages 106–110. Association for Computational Linguistics.
- Wang, W., Kennedy, R., Lazer, D., and Ramakrishnan, N. (2016). Growing pains for global monitoring of societal events. *Science*, 353(6307):1502–1503.
- Ward, M. D., Beger, A., Cutler, J., Dickenson, M., Dorff, C., and Radford, B. (2013). Comparing gdelt and icews event data. *Event Data Analysis*, 21(1):267–297.
- Weidmann, N. B. and Rød, E. G., (2019). *The Internet and Political Protest in Autocracies*, chapter Coding Protest Events in Autocracies. Oxford Studies in Digital Politics, Oxford.
- Yoruk, E. (2012). *The politics of the Turkish welfare system transformation in the neoliberal era: Welfare as mobilization and containment*. The Johns Hopkins University.