

Article

Not peer-reviewed version

Causal Extraction to Measure Media Bias? A Comparative Analysis of Causal Attributions in Reporting of the Israel-Palestine Conflict

[Hannah Béchara](#) ^{*}, Paulina Garcia Corral, Slava Jankin

Posted Date: 4 June 2024

doi: 10.20944/preprints202406.0189.v1

Keywords: Media Bias; Causal Relation Extraction; Machine Learning



Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Article

Causal Extraction to Measure Media Bias? A Comparative Analysis of Causal Attributions in Reporting of the Israel-Palestine Conflict

Hannah Béchara^{1,†,*} , Paulina Garcia Corral^{1,†} and Slava Jankin^{2,*}

¹ Data Science Lab – Hertie School; bechara@hertie-school.org, corral@hertie-school.org

² School of Government – University of Birmingham ; v.jankin@bham.ac.uk

* Correspondence: bechara@hertie-school.org

† These authors contributed equally to this work.

Abstract: Causal language can be a powerful tool for media outlets to subtly influence public perception and reinforce specific narratives. Therefore, identifying these linguistic patterns can be helpful in assessing the level of bias present in media content. In this paper, we directly compare the prevalence of causal language in headlines pertaining to the coverage of the Israel-Palestine conflict and the Russia-Ukraine war. We specifically look at the BBC and Al Jazeera headlines covering the assault on Gaza since October 7, 2023. We use UniCausal, a model for causal text mining, in order to automatically extract the cause and effect of each headline. Our findings reveal that BBC headlines tend to avoid directly attributing causality to Israel for the violence in Gaza, whereas Al Jazeera headlines explicitly identify Israel as the responsible party. This differential use of causal language suggests a potential bias in media reporting.

Keywords: media bias; causal relation extraction; machine learning

1. Introduction

Media bias, the subtle colouring of news reporting and information dissemination, has the potential to shape public opinion and influence political outcomes [1]. This bias can manifest in many forms, from selecting which stories to publish (coverage bias), to the framing of narratives (presentation bias), and even more subtle use of particular linguistic constructs [2]. Media bias is closely tied to disinformation, which characterises information based on truth, completeness, informativeness and deceptiveness [3].

This news slant is difficult to detect automatically, especially in cases where the slant comes from word choice or omission. This is due to the fact that slant is not necessarily reflected as false information, but rather stylistic choices that produce it. Additionally, for it to be considered bias, it must also be sustained across texts, and proven as such. Although there are multiple sources of news bias, we propose to study bias using a novel method by focusing on two specific sources:

- 1 **False cause fallacy** (*non causa pro causa*) [4]: This fallacy consists of mislocating the cause of one phenomenon to another that is only seemingly related. A media outlet may incorrectly causally link two events, such as attributing a rise in crime to a recent policy change without adequately exploring other factors that contributed.
- 2 **Missing cause bias** (information omission) [5]: a special type of omission that consistently omits attributing responsibility, placing blame or giving praise to specific acts or actors that caused an event, such as passively describing a violent attack by using sentences that do not contain a subject.

Although it may seem that measuring only two causes of bias limits the extent of bias assessment, causal relationships are highly relevant in news understanding and significantly shape opinions. Experiments in human cognition have shown that causal understanding is a primary cognitive task for comprehending the world and is influential in downstream tasks such as prediction and explanation [6]. Additionally, cause-effect understanding is intrinsically linked to temporal relationships and narrative

understanding [7], making it an information-rich and highly relevant semantic relationship that when extracted can be used to measure bias.

In this paper, we introduce a novel method to study news slant, by applying a Causal Extraction (CE) model to identify cause-effect pairs in news events. Causal relation extraction models classify and identify causal language in text [8]. We propose using causal relation extraction models to measure false cause attribution and missing cause in news media. By using this method, it allows us to quantify and compare assigned causes and omitted causes after matching documents in a news corpus of the same event. Since causal relations are a subset of semantic relations in text, focusing on these sources of bias reduces ambiguity by limiting researcher interpretation of what the relationship between the extracted spans is. This approach provides a new measure for bias that extends beyond labour-intensive supervised classification, focusing instead on one single relevant semantic relationship.

2. Related Work

2.1. Bias Detection

While Media Bias has a long tradition in the realms of social sciences, it remains a relatively young research topic in Natural Language Processing. Common NLP techniques such as sentiment analysis [9], topic modelling [10] and lexical feature analysis [11] have been used to detect the presence of media bias. More recently, supervised machine learning classifiers trained on transformer-based models have achieved better results although some underlying issues persist. Altogether, media bias is a multi-faceted problem with several competing definitions, and bias detection remains a complex task.

Several researchers have addressed the detection of media bias as a classification problem. This classification can be binary (either bias exists or it does not exist), or it can be treated as a multi-classification problem. One such paper, which rates articles by degree of polarisation, identifies a set of 130 content-based features that span 7 categories: structure, complexity, sentiment, bias, morality, topic and engagement [12], and show that they all contribute towards media bias detection. In contrast, media bias has also been classified by its stance towards an event [13] or by its place on the political compass [14].

Given that supervised classification usually uses case-specific annotations, there are some relevant attempts for our case study. To detect media bias using NLP, [15] compared three supervised machine learning algorithms trained on a Israel - Palestine conflict news dataset, a SVM with bio-grams achieved the highest performance of 91.76% accuracy and F-score of 91.46%. [16] collected data from history book excerpts and newspaper articles of the same conflict, and trained a sequence classifiers to predict authorship provenance. Their best model to detect narrative origin achieved an F1 score of 85.1% for history book excerpts and 91.9% for newspaper articles. Additionally, [17] manually classified pro-Russia and pro-Western bias of news articles using a baseline SVM classifier with doc2vec embeddings, that achieved an F1 score of 86%. However, their results suggest that models may be learning journalistic styles rather than actually modelling bias. [18] applied a novel methodology to gold-labelled set of articles annotated for Pro-Russian bias, where the a Naive Bayes classifier achieved 82.6% accuracy and a feed-forward neural networks achieved 85.6% accuracy.

2.2. Causal Language Modelling

Causality mining is the task of identifying causal language that connects events in a text, it is “an information extraction step that acquires causal relations from a collection of documents” [8]. Causal language can be expressed explicitly or implicitly. The former tends to use connectors such as “because” or “therefore” to signal a causal relationship between events; causal verbs can also be used to express causality. Causality can be found inter or intra-sentential. Because causal language is a linguistically, syntactically and semantically complex construction used to express causal reasoning [19,20], it tends

to rely on contextual information, and yields low inter-annotator agreement when annotated by humans [21].

Causal language modelling is usually broken into three steps: 1) causal sequence classification, 2) causal extraction and 3) pair classification [22]. It's downstream applications include biomedical discovery, emergency management, and financial narrative summaries. Causal mining used to be based on pattern matching by identifying causal connectors [8]. Advances in Machine Learning allowed for statistical pattern recognition models, such as SVMs and Bayesian models [23,24]. With the introduction of deep learning, a combination of architectures that include CNNs [25,26], CRFs [27], or LSTMs [28,29] improved previous results. Moreover, fine-tuning transformer-based models such as BERT significantly improved both classification and extraction capabilities in different genres [22,30,31]. Despite the vast improvement LLMs show in most NLP tasks, there is a lack of research on Causal language modelling. Notably, Shukla *et al.* [32] evaluated causality detection in financial texts for English and Spanish using GPT models, yielding high results using few-shot learning.

2.3. Applications of Causal Language Modelling in Media Analysis

[33] tested the hypothesis that more causally and semantically coherent text documents are more likely to be shared by individuals on social media. Results show that the role of causally and semantically coherent content in promoting online sharing and motivate better measures of these key constructs. FinCausal [34] analyses narratives in financial news articles, detecting if something is considered a cause for a prior event. Therefore, it aims to develop an ability to explain the the cause of a change in the financial landscape, as a preamble to generating financial narrative summaries. Finally, [Anonymous] created a dataset and benchmarked it for causal language detection in political press conferences, to better analyse political discourse quality.

3. Experimental Setup

Our bias detection method using causal extraction is composed of four steps: First, we use a causal sequence classification model on the headlines from the news we collected and matched. Second, we will extract the cause-effect pairs using a causal extraction model on the positive sequences. Third, we will compare cause selection as well as identify missing cause on the matched positive pairs. To demonstrate that the differences are not due to publishing stylistic differences, we repeated the same process with data from a different conflict, the Russia-Ukraine war. Finally, we run the previous steps and compare each media's use of causality to itself in a different conflict.

3.1. Dataset Collection

We gathered headlines from Al Jazeera English and the BBC pertaining to the Israel-Palestine conflict, dating back to October 7, 2023. We also scrapped headlines from the same media from the Ukraine - Russia war, from Al Jazeera English is an English-language news channel headquartered in the Middle East and funded in part by the Qatari government. The BBC (British Broadcasting Corporation) is a British public service broadcaster, the oldest and largest in the United Kingdom, and is funded principally by a licence fee charged by the British Government. As such, both outlets can be considered "state media" and mainstream of their respective governments, for the purposes of this study. After filtering out articles with no relevance to either of the conflicts, the total number of articles are summarised in Table 1. All the data is available in our repository¹. We matched up the headlines using cosine similarity and a time window of 2 days, resulting in a side-to-side comparison of each outlet's coverage of the same event. We consider values above 0.5 in the distance measure, and within two days of each other, to be same event matches, summarised in Table 2. Table 3 shows a sample of these aligned headlines.

¹ www.github.com/our_repository

Table 1. Articles collected from the BBC and Al Jazeera for Gaza conflict.

Source	Number of Articles
BBC Middle East ²	793
AlJazeera Israel-Palestine Conflict ³	1465
BBC Ukraine	784
AlJazeera Ukraine - Russia War	1018

Table 2. Alligned Articles Dataset.

Conflict	Number of Articles
Israel - Palestine	98
Ukraine - Russia	65

Table 3. A Sample of Aligned Headlines from BBC and AlJazeera.

BBC	Al Jazeera
Israel - Palestine Conflict Headlines	
Refaat Alareer: Palestinians mourn writer killed in air strike	Palestinians mourn poet Refaat Alareer killed in Israeli air strike
ICJ says Israel must prevent genocide in Gaza	ICJ orders Israel to prevent acts of genocide in Gaza
Hamas deputy leader Saleh al-Arouri killed in Beirut blast	Senior Hamas official Saleh al-Arouri killed in Beirut suburb
Ukraine - Russia Conflict Headlines	
Rustem Umerov: Who is Ukraine’s next defence minister?	Who is Rustem Umerov, Ukraine’s next defence minister?
Ukraine and Russia complete first prisoner swap since plane crash	Russia and Ukraine complete first prisoner exchange since plane crash
Ukraine celebrates first Christmas on 25 December	Ukraine officially celebrates Christmas on December 25 for the first time

3.2. Causal Extraction Models to Detect Media Bias

We classified and annotated the collected headlines (c.f. Section 3.1) using an open source and freely available repository UniCausal⁴. While we collected the full articles, we focus only on the headlines for this study, as they are specifically crafted to appeal to the reader and capture their attention. UniCausal is built off a BERT [35] pre-trained language model fine-tuned on six, high-quality human-annotated corpora. UniCausal is especially well-suited for our task as five of these six causal corpora are News corpora, replicated in Table 4, as reported by the authors.

Table 4. The six causal corpora used to fine-tune the BERT Model.

Corpus	Source
AltLex [36]	News
BECAUSE 2.0 [37]	News, Congress Hearings
CausalTimeBank (CTB) [38]	News
EventStoryLine V1.0 (ESL) [7]	News
Penn Discourse Treebank V3.0 (PDTB) [39]	News
SemEval 2010 Task 8 (SemEval) [40]	Web

UniCausal [22] offers a unified benchmark and resource for causal text mining standardised across three tasks:

- Causal Sequence Classification: Returns a causal label (either *causal* or *non-causal*)
- Cause-Effect Span Detection: Identifies the text related to the Cause-Effect spans, i.e., which words correspond to the *cause* and *effect* arguments.
- Causal Pair Classification: Identifies whether a marked entity pair is causally related.

For the purposes of our study, we focus on the first two tasks: Causal Sequence Classification, and Cause-Effect Span Detection. Causal Sequence Classification returns a label of 1 (causal) or 0 (non-causal), in addition to a confidence score for each label. Figure 1 shows a sample sentence from the headlines, with LABEL_1 representing a causal score.

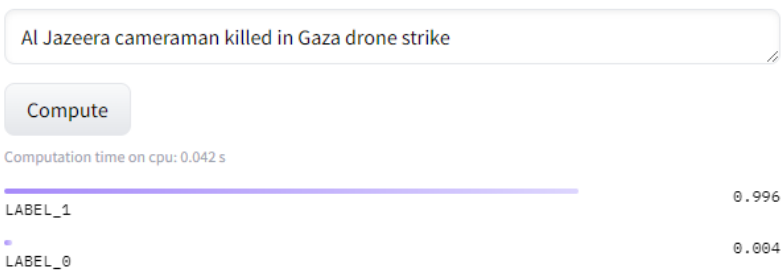


Figure 1. A Sample Annotation by UniCausal using the Sequence Classifier

The Cause-Effect Span Detection task is a text mining task that, given a causal example, identifies which words correspond to the “cause”, and which words correspond to the “effect”. UniCausal’s Span Detection task can identify up to 3 causal relations and their spans.

⁴ <https://github.com/tanfiona/UniCausal>

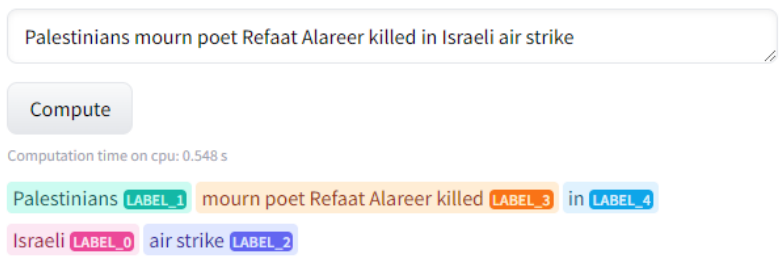


Figure 2. A Sample Annotation by UniCausal using Span Detection

We run both tasks on all the collected headlines to determine which sentences are causal, and what the identified explicit cause is presented in each of the headlines. We do not perform any Causal Pair Classification for this study.

4. Results

This section summarises our quantitative results from the experiments described in Section 3. We have divided it into 2 subsections. Section 4.1 compares the BBC and AlJazeera coverage of the Israel-Palestine conflict since October 7, 2023. Section 4.2 compares BBC’s Russia-Ukraine coverage to the same outlet’s Israel-Palestine coverage.

4.1. BBC vs AlJazeera Coverage

We compare the results of the Causal Sequence Classifier run on the “BBC Middle East” headlines vs the “AlJazeera Israel-Palestine” headlines. We find that UniCausal tags more BBC headlines as causal, with 28.4% of headlines having a causal construction. In contrast, only 24% of AlJazeera’s relevant headlines include a causal relation. As this difference seems to be within the margin of error of UniCausal, there does not appear to be a discrepancy between the causal constructions in the BBC vs AlJazeera headlines. The Cause-Effect Span detection task shows a different story, however. When we search the causal headlines for terms like “Israel”, “IDF” and “Israeli” in the “Cause” span, we find these terms in only 20 BBC headlines (0.08% of all headlines with causal constructions). This compares to 206 headlines in the AlJazeera headlines (57% of all causal sentences). Furthermore, we select for sentences in which the “effect” span includes references to “killing” or “death”. Of these sentences, only 1% of BBC headlines also include terms related to “Israel” or the “IDF” in the causal span as opposed to 54 out of 112 headlines (54%) of AlJazeera headlines.

Table 5 summarises these results.

Table 5. BBC vs AlJazeera Headline Breakdown of Causal Sentences.

Headlines	BBC Middle East	Al Jazeera	BBC Russia-Ukraine
Number of Causal Sentences			
Total Headlines	793	1465	784
Causal Headlines	225 (28.4%)	350 (24%)	236 (30%)
Causal Sentences with cause spans that include terms related to Israel/Russia			
Number of Headlines	20 (8%)	206 (57%)	98 (41.5%)
Causal Sentences with effect spans that include terms related to killed or dead			
Number of Headlines	78	112	47
Combined	8 (1%)	54 (48%)	16 (34%)

Furthermore, we study the 89 headlines from BBC and AlJazeera that we aligned using cosine similarity. This allows us to directly compare coverage of the same events from two disparate sources. The results show a similar breakdown as in Table 4, with only 15 out of 87 headlines with causal constructions. Because of the small number of headlines, we are able to manually check each of these

aligned texts and compare them directly for more insight into the results prentenced in Table 5. Here we find that only one of the BBC headlines explicitly mentions “Israel” in the cause of airstrikes in the headline. Table 6 shows some of these headlines side-by side, highlighting the slant in the BBC vs AlJazeera. In these examples, all the sentences are identified as causal, and the “cause” span is in bold.

Table 6. A direct comparison of headlines with causal constructions – BBC vs Aljazeera

BBC	Al Jazeera
Israel-Palestine Conflict Headlines	
Samer Abudaqa: Al Jazeera cameraman killed in Gaza drone strike	Al Jazeera journalist Samer Abudaqa killed in Israeli attack in Gaza
Wael Al-Dahdouh: Al Jazeera reporter’s family killed in Gaza strike	Al Jazeera condemns Israeli killing of journalist Wael Al-Dahdouh’s family
Al Jazeera bureau chief’s son Hamza al-Dahdouh among journalists killed in Gaza	Hamza son of Al Jazeera’s Wael Dahdouh killed in Israeli attack in Gaza

4.2. BBC's Coverage of the Russia-Ukraine Conflict

We further compare BBC's coverage of the Russia-Ukraine conflict to its coverage of the Israel-Palestine conflict. The intuition behind this additional set of experiments is to explore whether or not the reluctance to mention the cause in the headlines is specific to the Israel-Palestine conflict, or if it is simply a feature of BBC headlines. The former might indicate a bias in reporting. UniCausal's Causal Sequence Classifier finds that about 236 (30%) of the 784 headlines on the Russia-Ukraine conflict contain a causal relationship, as compared to only 28% of the Israel-Palestine headlines. The Cause-Effect Span detection shows that BBC is much more likely to identify Russia as the cause in headlines, with terms like "Russia" and "Russian" present in 41.5% of the causal spans. We can conclude from these results (summarised in Table 5), that the reluctance to name the perpetrator in the headlines is unique to the Israel-Palestine BBC coverage, and therefore a result of slanted reporting.

5. Discussion

The framing of conflict-related events in media reporting can significantly influence public perception. This is particularly evident in the coverage of the Israel-Palestine conflict, where the attribution of causality in incidents of death and destruction is a contentious issue. Our comparative analysis of BBC and Al Jazeera articles reveals a discernible difference in how each media outlet attributes causality, which may indicate underlying biases.

BBC articles on the conflict tend to present information in a more neutral or detached manner, often avoiding direct attribution of causality to Israel for the deaths and destruction in Gaza. For example, phrases like "reported killed in latest strikes" or "scores were killed in the camp" are used without explicitly identifying Israel as the cause. This is evident in that only 1% of causal sentences that describe death and destruction will attribute the cause directly to Israel. This tendency reflects an implicit bias through omission, thereby absolving one party from direct responsibility.

In contrast, Al Jazeera headlines are more forthright in attributing causality, with statements such as "Israeli tanks in Rafah city as new deadly strike targets Gaza civilians" and "Israel has attacked twice since Sunday, killing tens of people in horrific circumstance" clearly identifying Israel as the agent of action. This is reflected in our results, where a more balanced 57% of headlines that include death and destruction include a direct mention of Israel as the cause. This direct attribution of responsibility aligns with Al Jazeera's editorial stance, which is often perceived as being more sympathetic to the Palestinian perspective.

The variance in causal language between these two outlets highlights the role of media framing in conflict reporting. While BBC's approach could be seen as an effort to remain impartial, it may also suggest a reluctance to assign blame to Israel, potentially underplaying the severity of the situation in Gaza. On the other hand, Al Jazeera's explicit attribution of causality to Israel reflects a different editorial choice, one that emphasises the impact of Israeli actions on Palestinian civilians.

6. Conclusions

In this paper, we explored the use of causal language in media reporting and how its detection can act as an indicator of bias, offering a window into the subtle ways in which narratives are shaped. We compared headlines from two different media outlets, the BBC and AlJazeera, pertaining to their reporting on two different conflicts: The Israel-Palestine conflict, and the Russia-Ukraine conflict. Using state-of-the-art causal extraction methods, we automatically classify the headlines as causal and non-causal. We further extract the cause and effect spans of each of the headlines. A comparison shows a clear bias by omission on the part of the BBC Israel-Palestine reporting as opposed to AlJazeera's reporting, and even BBC's own reporting on the Russia-Ukraine conflict.

Our research is not without its limitations. The scope of the study was confined to just two media outlets, which do not represent the entire spectrum of journalistic practices. Further research could expand upon this work and incorporate headlines from different sources, including different languages and from various political leanings. Furthermore, this study focuses on headlines only, as they are

crafted to capture the most attention. However, a future avenue of research could also focus on the articles themselves and the causal language and slant present therein.

Despite the limitation, this research has demonstrated that the presence and patterns of causative constructions can reflect an outlet's underlying biases, whether intentional or not. Our findings suggest that the detection of causal language can serve as a powerful tool for media analysts and consumers alike, promoting a more discerning and critical approach to news consumption. While no single method can encapsulate the full spectrum of media bias, the identification of causal language provides a measurable and scalable means to gauge the impartiality of news content.

Author Contributions: All three authors contributed to the conceptualisation and methodology of this project. Hannah Béchara and Paulina Garcia Corral contributed formal analysis, writing – original draft preparation, review and editing and data curation. Slava Jankin contributed to funding acquisition, supervision and project administration.

Funding: This project has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101057131, Climate Action To Advance HeaLthY Societies in Europe (CATALYSE), and from the DFG (EXC number 2055 – Project number 390715649, SCRIPTS)

Data Availability Statement: The original data presented in the study are openly available in BBCProject at <https://github.com/hbechara/BBCProject>.

7. Acknowledgements

The authors thank the DFG (EXC number 2055 – Project number 390715649, SCRIPTS) for funding a part of this research. This project has also received funding from the European Union's Horizon Europe research and innovation program under Grant Agreement No 101057131, Climate Action To Advance HeaLthY Societies in Europe (CATALYSE).

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

NLP	Natural Language Processing
CE	Causal Extraction
SVM	Support Vector Machines
CNN	Convolutional Neural Network
CRF	Conditional Random Fields
LSTM	Long Short-Term Memory
LLM	Large Language Model
GPT	Generative Pre-Trained Transformer
BERT	Bidirectional Encoder Representations from Transformers

References

1. Bernhardt, D.; Krasa, S.; Polborn, M. Political polarization and the electoral effects of media bias. *Journal of Public Economics* **2008**, *92*, 1092–1104. doi:<https://doi.org/10.1016/j.jpubeco.2008.01.006>.
2. Saez-Trumper, D.; Castillo, C.; Lalmas, M. Social media news communities: Gatekeeping, coverage, and statement bias. *Proceedings of the 22nd ACM international conference on Information & Knowledge Management*, 2013, pp. 1679–1684.
3. Karlova, N.; Fisher, K.E. A social diffusion model of misinformation and disinformation for understanding human information behaviour. *Inf. Res.* **2013**, *18*.
4. Manninen, B.A. False Cause. In *Bad Arguments*; Arp, R.; Barbone, S.; Bruce, M., Eds.; Wiley, 2018. doi:10.1002/9781119165811.ch79.
5. Gentzkow, M.; Shapiro, J.M. Media Bias and Reputation. *Journal of Political Economy* **2006**, *114*, 280–316. doi:10.1086/499414.

6. Waldmann, M.; Hagmayer, Y. Causal reasoning. In *Oxford Handbook of Cognitive Psychology*; Oxford University Press, 2013; pp. 733–752.
7. Caselli, T.; Vossen, P. The Event StoryLine Corpus: A New Benchmark for Causal and Temporal Relation Extraction. Proceedings of the Events and Stories in the News Workshop; Caselli, T.; Miller, B.; van Erp, M.; Vossen, P.; Palmer, M.; Hovy, E.; Mitamura, T.; Caswell, D., Eds.; Association for Computational Linguistics: Vancouver, Canada, 2017; pp. 77–86. doi:10.18653/v1/W17-2711.
8. Drury, B.; Gonçalves Oliveira, H.; de Andrade Lopes, A. A survey of the extraction and applications of causal relations. *Natural Language Engineering* **2022**. doi:10.1017/S135132492100036X.
9. Lin, Y.R.; Bagrow, J.; Lazer, D. More voices than ever? quantifying media bias in networks. Proceedings of the international AAAI conference on web and social media, 2011, Vol. 5, pp. 193–200.
10. Best, C.; van der Goot, E.; Blackler, K.; Garcia, T.; Horby, D. Europe media monitor. *Technical Report EUR221 73 EN, European Commission* **2005**.
11. Hube, C.; Fetahu, B. Detecting biased statements in wikipedia. Companion proceedings of the the web conference 2018, 2018, pp. 1779–1786.
12. Horne, B.; Khedr, S.; Adali, S. Sampling the news producers: A large news and feature data set for the study of the complex media landscape. Proceedings of the International AAAI Conference on Web and Social Media, 2018, Vol. 12.
13. Cremisini, A.; Aguilar, D.; Finlayson, M.A. A challenging dataset for bias detection: The case of the crisis in the ukraine. Social, Cultural, and Behavioral Modeling: 12th International Conference, SBP-BRIMS 2019, Washington, DC, USA, July 9–12, 2019, Proceedings 12. Springer, 2019, pp. 173–183.
14. Baly, R.; Karadzhov, G.; An, J.; Kwak, H.; Dinkov, Y.; Ali, A.; Glass, J.; Nakov, P. What was written vs. who read it: News media profiling using text analysis and social media context. *arXiv preprint arXiv:2005.04518* **2020**.
15. Al-Sarraj, W.F.; Lubbad, H.M. Bias Detection of Palestinian/Israeli Conflict in Western Media: A Sentiment Analysis Experimental Study. 2018 International Conference on Promising Electronic Technologies (ICPET), 2018, pp. 98–103. doi:10.1109/ICPET.2018.00024.
16. Wei, J.; Santos, E. Narrative Origin Classification of Israeli-Palestinian Conflict Texts. The Thirty-Third International FLAIRS Conference (FLAIRS-33), 2020.
17. Cremisini, A.; Aguilar, D.; Finlayson, M.A. A Challenging Dataset for Bias Detection: The Case of the Crisis in the Ukraine. Social, Cultural, and Behavioral Modeling. SBP-BRIMS 2019; Thomson, R.; Bisgin, H.; Dancy, C.; Hyder, A., Eds. Springer, Cham, 2019, Vol. 11549, *Lecture Notes in Computer Science*. doi:10.1007/978-3-030-21741-9_18.
18. Potash, P.; Romanov, A.; Gronas, M.; Rumshisky, A.; Gronas, M. Tracking Bias in News Sources Using Social Media: The Russia-Ukraine Maidan Crisis of 2013–2014. Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism; Popescu, O.; Strapparava, C., Eds.; Association for Computational Linguistics: Copenhagen, Denmark, 2017; pp. 13–18. doi:10.18653/v1/W17-4203.
19. Solstad, T.; Bott, O. 619Causality and Causal Reasoning in Natural Language. In *The Oxford Handbook of Causal Reasoning*; Oxford University Press, 2017; https://doi.org/10.1093/oxfordhb/9780199399550.013.32.
20. Neeleman, A.; Van de Koot, H. The Linguistic Expression of Causation. In *The Theta System: Argument Structure at the Interface*; Oxford University Press: Oxford, 2012. https://doi.org/https://doi.org/10.1093/acprof:oso/9780199602513.003.0002.
21. Dunietz, J.; Levin, L.; Carbonell, J. The BECauSE Corpus 2.0: Annotating Causality and Overlapping Relations. Proceedings of the 11th Linguistic Annotation Workshop; Association for Computational Linguistics: Valencia, Spain, 2017; pp. 95–104.
22. Tan, F.A.; Zuo, X.; Ng, S. UniCausal: Unified Benchmark and Repository for Causal Text Mining. Big Data Analytics and Knowledge Discovery - 25th International Conference, DaWaK 2023, Penang, Malaysia, August 28–30, 2023, Proceedings; Wrembel, R.; Gamper, J.; Kotsis, G.; Tjoa, A.M.; Khalil, I., Eds. Springer, 2023, Vol. 14148, *Lecture Notes in Computer Science*, pp. 248–262. doi:10.1007/978-3-031-39831-5_23.
23. Hidey, C.; Mckeown, K. Identifying Causal Relations Using Parallel Wikipedia Articles. *54th Annual Meeting of the Association for Computational Linguistics, ACL 2016 - Long Papers* **2016**, 3, 1424–1433. Publisher: Association for Computational Linguistics (ACL), doi:10.18653/V1/P16-1135.

24. Zhao, S.; Wang, Q.; Massung, S.; Qin, B.; Liu, T.; Wang, B.; Zhai, C. Constructing and Embedding Abstract Event Causality Networks from Text Snippets. *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*. ACM, 2017, pp. 335–344. doi:10.1145/3018661.3018707.
25. Kruengkrai, C.; Torisawa, K.; Hashimoto, C.; Kloetzer, J.; Oh, J.H.; Tanaka, M. Improving Event Causality Recognition with Multiple Background Knowledge Sources Using Multi-Column Convolutional Neural Networks. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2017, Vol. 31. <https://doi.org/10.1609/aaai.v31i1.11005>.
26. de Silva, T.N.; Zhibo, X.; Rui, Z.; Kezhi, M. Causal Relation Identification Using Convolutional Neural Networks and Knowledge Based Features. *International Journal of Computer and Systems Engineering* **2017**, *11*.
27. Fu, J.F.; Liu, Z.T.; Liu, W.; Zhou, W. Event Causal Relation Extraction Based on Cascaded Conditional Random Fields. *Pattern Recognition and Artificial Intelligence* **2011**, *24*, 567.
28. Li, Z.; Li, Q.; Zou, X.; Ren, J. Causality Extraction based on Self-Attentive BiLSTM-CRF with Transferred Embeddings. *ArXiv* **2019**, *abs/1904.07629*.
29. Dasgupta, T.; Saha, R.; Dey, L.; Naskar, A. Automatic Extraction of Causal Relations from Text using Linguistically Informed Deep Neural Networks. *Proceedings of the SIGDIAL 2018 Conference; Association for Computational Linguistics: Melbourne, Australia*, 2018; pp. 12–14.
30. Yang, J.; Xiong, H.; Zhang, H.; Hu, M.; An, N. Causal Pattern Representation Learning for Extracting Causality from Literature. *Proceedings of the 2022 5th International Conference on Machine Learning and Natural Language Processing; Association for Computing Machinery: New York, NY, USA*, 2023; MLNLP '22, p. 229–233. doi:10.1145/3578741.3578787.
31. Khetan, V.; Ramnani, R.R.; Anand, M.; Sengupta, S.; Fano, A.E. Causal BERT: Language Models for Causality Detection Between Events Expressed in Text. *Intelligent Computing. Lecture Notes in Networks and Systems; Arai, K., Ed., 2022, Vol. vol 283*. doi:https://doi.org/10.1007/978-3-030-80119-9_64.
32. Shukla, N.K.; Katikeri, R.; Raja, M.; Sivam, G.; Yadav, S.; Vaid, A.; Prabhakararao, S. Investigating Large Language Models for Financial Causality Detection in Multilingual Setup. *2023 IEEE International Conference on Big Data (BigData)*, 2023, pp. 2866–2871. doi:10.1109/BigData59044.2023.10386558.
33. Hosseini, P.; Diab, M.; Broniatowski, D.A. Does Causal Coherence Predict Online Spread of Social Media? Social, Cultural, and Behavioral Modeling. *SBP-BRIMS 2019; Thomson, R.; Bisgin, H.; Dancy, C.; Hyder, A., Eds. Springer, Cham*, 2019, Vol. 11549, *Lecture Notes in Computer Science*. doi:10.1007/978-3-030-21741-9_19.
34. Mariko, D.; Abi-Akl, H.; Trottier, K.; El-Haj, M. The Financial Causality Extraction Shared Task (FinCausal 2022). *Proceedings of the 4th Financial Narrative Processing Workshop @LREC2022; El-Haj, M.; Rayson, P.; Zmandar, N., Eds.; European Language Resources Association: Marseille, France*, 2022; pp. 105–107.
35. Devlin, J.; Chang, M.; Lee, K.; Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR* **2018**, *abs/1810.04805*, [[1810.04805](https://arxiv.org/abs/1810.04805)].
36. Hidey, C.; McKeown, K. Identifying Causal Relations Using Parallel Wikipedia Articles. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); Erk, K.; Smith, N.A., Eds.; Association for Computational Linguistics: Berlin, Germany*, 2016; pp. 1424–1433. doi:10.18653/v1/P16-1135.
37. Dunietz, J.; Levin, L.; Carbonell, J. The BECauSE Corpus 2.0: Annotating Causality and Overlapping Relations. *Proceedings of the 11th Linguistic Annotation Workshop; Schneider, N.; Xue, N., Eds.; Association for Computational Linguistics: Valencia, Spain*, 2017; pp. 95–104. doi:10.18653/v1/W17-0812.
38. Mirza, P.; Sprugnoli, R.; Tonelli, S.; Speranza, M. Annotating Causality in the TempEval-3 Corpus. *Proceedings of the EACL 2014 Workshop on Computational Approaches to Causality in Language (CAtoCL); Kolomiyets, O.; Moens, M.F.; Palmer, M.; Pustejovsky, J.; Bethard, S., Eds.; Association for Computational Linguistics: Gothenburg, Sweden*, 2014; pp. 10–19. doi:10.3115/v1/W14-0702.

39. Webber, B.; Prasad, R.; Lee, A.; Joshi, A. The penn discourse treebank 3.0 annotation manual. *Philadelphia, University of Pennsylvania* **2019**, 35, 108.
40. Hendrickx, I.; Kim, S.N.; Kozareva, Z.; Nakov, P.; Ó Séaghdha, D.; Padó, S.; Pennacchiotti, M.; Romano, L.; Szpakowicz, S. SemEval-2010 Task 8: Multi-Way Classification of Semantic Relations between Pairs of Nominals. *Proceedings of the 5th International Workshop on Semantic Evaluation*; Erk, K.; Strapparava, C., Eds.; Association for Computational Linguistics: Uppsala, Sweden, 2010; pp. 33–38.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.