

Overview of the CLEF-2023 CheckThat! Lab Task 3 on Political Bias of News Articles and News Media

Giovanni Da San Martino^{1,*}, Firoj Alam^{2,*}, Maram Hasanain¹, Rabindra Nath Nandi³, Dilshod Azizov⁴ and Preslav Nakov⁴

¹University of Padova, Italy

²Qatar Computing Research Institute, HBKU, Qatar

³Hishab Singapore Pte. Ltd, Singapore

⁴Mohamed bin Zayed University of Artificial Intelligence, UAE

Abstract

We provide an overview of task 3 of the CheckThat! lab from the Cross-Language Evaluation Forum (CLEF) 2023, which focuses on predicting the political leaning of English-language news articles and news media outlets. We describe the data collection, the task setup, the evaluation outcomes, and the approaches used by the participating teams. A total of six teams submitted runs for the two subtasks. The top-performing system in Subtask 3A achieved a Mean Absolute Error (MAE) of 0.473, while the best system in Subtask 3B yielded a MAE of 0.549. We make all datasets and evaluation scripts available to the public, aiming to boost further research on this problem.

Keywords

Political bias, news articles, news media.

1. Introduction

In the era of widespread digital information, the impact of political bias in news media and news articles has emerged as a significant concern for democratic societies [1]. Accusations of bias against news organizations, which could influence the public opinion and the policy discourse, have been longstanding [2]. Various research approaches have been used to detect the potential bias of news articles, e.g., by using a headline attention network [3] or by monitoring the frequency of mentions and quotes of politicians from different political parties [4]. At the medium level, a number of recent studies [5, 6, 7, 8, 9, 8, 10, 11, 12] have been conducted, making use of variety of information sources ranging from news articles to tweets, YouTube channels, and user overlap.

To enhance social awareness and to counteract the spread of false information, the CheckThat! lab at CLEF [13, 14, 15, 16, 17, 18, 19] has developed tasks using high-quality data and suitable evaluation measures. As part of this initiative, the CheckThat! lab organizers have offered five distinct tasks [20, 21, 22].

CLEF 2023: Conference and Labs of the Evaluation Forum, September 18–21, 2023, Thessaloniki, Greece

*Corresponding author.

✉ dasan@math.unipd.it (G. Da San Martino); fialam@hbku.edu.qa (F. Alam); mhasanain@hbku.edu.qa (M. Hasanain); rabindro.rath@gmail.com (R. N. Nandi); dilshod.azizov@mbzuai.ac.ae (D. Azizov); preslav.nakov@mbzuai.ac.ae (P. Nakov)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings (CEUR-WS.org)

In this paper, we provide an overview of Task 3, which focuses on detecting the political leaning of *news articles* and *media outlets*. We release high-quality, manually annotated data for these two subtasks, based on a 3-point ordinal scale for modeling the bias at the article and at the medium level.

The remainder of this paper is organized as follows: In Section 2, we outline the task at hand and describe the dataset we created and released. Section 3 offers a comprehensive overview of the evaluation settings. We discuss the results and we delve into the details of the submitted systems in Section 4. Section 5 brings to light previous and recent work that aligns with our study. Finally, Section 6 presents our concluding observations and suggests directions for future research.

2. Task and Datasets

Below, we define the tasks, and then we discuss the datasets.

2.1. Task definition

The goal of the task is to detect the political bias of news reporting at the article and at the media level. This is an ordinal classification task and it is offered in English. It includes two subtasks, defined below.

Subtask 3A: Political Bias of News Articles Given an article, classify its political leaning as left, center, or right.

Subtask 3B: Political Bias of News Media Given the news article(s) a news outlet (e.g., www.cnn.com), predict the overall political bias of that news outlet as left, center, or right.

2.2. Datasets

Below, we describe the datasets for the two subtasks.

2.2.1. Subtask 3A


We release a new dataset, which we crawled from AllSides¹, a website that gathers news articles from a variety of reputable national and international news sources to ensure a balanced representation across different political spectrums. The site offers meticulous annotations of the bias of news articles, including expert assessments, third-party analysis, independent evaluations, and community input. In addition, AllSides uses annotated articles to support its Balanced Search tool, which displays the news coverage of a specific issue from numerous media providers, each with a different political bias from all political perspectives, as depicted in Figure 1.

To ensure that the dataset remains relevant and reflects the current political environment, it includes news articles published from late 2022 till early 2023. Each article has several attributes, as shown in Table 1. In total, we have just over 55k articles in the dataset. We provide statistics about the dataset in Table 2.

¹www.allsides.com

Figure 1: Examples of news articles with bias labels assigned by AllSides (Subtask 3A). Source: www.allsides.com

Balanced News from the **Left**, **Center** and **Right**



HEADLINE ROUNDUP ?

Inflation Rose 4% Annually in May, Lowest in 2 Years

The Consumer Price Index (CPI) rose 4% for the year ending in May, the lowest annual inflation rate since March 2021, according to the Bureau of Labor Statistics. The Details: Prices rose 0...

Robert Nickelsberg / Getty Images file

From the Left

Inflation cooled to 4% in May, the lowest reading since March 2021

NBC News (Online) L L C R R

From the Center

Inflation rose at a 4% annual rate in May, the lowest in 2 years

CNBC L L C R R

From the Right

Inflation cools sharply in May to 4%, lowest in 2 years

Fox Business L L C R R

Table 1
Data attributes for the article (Subtask 3A).

ID	Unique identifier
Title	Headline
Content	Full text of the article
Label	Political bias of the article: left, center, or right

Table 2
Statistics about the training, development, and test partitions (Subtask 3A).

	Left	Center	Right	Total
Train	12,073	15,449	17,544	45,066
Dev	1,342	1,717	1,949	5,008
Test	2,589	1,959	650	5,198
Total	16,004	19,125	20,143	55,272

2.2.2. Subtask 3B

We assess the political bias of English-language media, sourced from Media Bias/Fact Check website.² On that website, experts conduct an in-depth analysis and annotate the political bias of entire news outlets: examples are shown in Table 3. We further include a certain number of articles, which we crawl from each media source: these are to be used by the participants to analyze that source. The dataset has similar attributes to subtask 3A, plus the source (the name of the medium) as an additional attribute. We have over 8,000 articles (approximately 10 per source) and over 1,000 news sources. Tables 4 and 5 show the label distribution and the number of articles and news media.

²www.mediabiasfactcheck.com

Table 3

Examples of news outlets and their biases (Subtask 3B).

Left	Center	Right
The Guardian	BBC News	Fox News
The New York Times	Reuters	The Daily Caller
The Washington Post	The Associated Press	The National Review

Table 4

Statistics about the training, the development, and the test partitions for the news outlets (Subtask 3B).

	Left	Center	Right	Total
Train	216	296	305	817
Dev	31	34	39	104
Test	25	29	48	102
Total	272	359	392	1,023

Table 5

Statistics about the training, the development, and the test partitions for news articles across all news outlets (Subtask 3B).

	Left	Center	Right	Total
Train	1,350	1,822	2,051	5,223
Dev	378	386	434	1,196
Test	526	536	564	1,626
Total	2,254	2,744	3,049	8,047

3. Evaluation Settings

Settings The evaluation comprises development and test phases. During the development phase, we provided the participants with the training and the development sets. This enabled them to internally validate their systems and to adjust the parameter values using the development set. During the test phase, the participants submitted their system’s predictions for the provided test set, which did not include reference labels. They were allowed to submit as many runs as they wanted, but only the last submission was considered as the final one.

Evaluation This is an ordinal classification task, and thus we used mean absolute error as the official measure for both subtasks.

Table 6

Results on the leaderboard: political bias of news articles and news media (MAE score).

Subtask 3A			Subtask 3B		
Rank	Team	MAE	Rank	Team	MAE
1	Accenture [24]	0.473	1	Accenture [24]	0.549
2	TOBB ETU [25]	0.646	2	Awakened	0.765
3	KUCST	0.736	3	Baseline	0.902
4	Awakened	0.752			
5	Baseline	0.877			
	Frank [23]	0.270		Frank [23]	0.320

4. Results and Overview of the Systems

4.1. Results

Table 6, shows the results for Task 3, in which four official teams and one non-official team participated. All teams outperformed the baseline.

Four teams participated in Subtask 3A, with Accenture taking the lead, having the lowest MAE of 0.473. They are followed by *TOBB ETU*, *KUCST*, and *Awakened* with MAE scores of 0.646, 0.736, and 0.752, respectively.

Two teams took part in Subtask 3B. Once again, Accenture was first, with a MAE of 0.549, followed by *Awakened*, with a MAE value of 0.765.

Team Frank [23], which did not officially appear on the leaderboard, outperformed all other teams on both subtasks. They achieved a MAE of 0.270 for subtask 3A and 0.320 in subtask 3B. Both scores are much lower than those of the participating teams.

4.2. Overview of the Systems

Accenture [24] used machine back-translation to augment the minority classes examples and thus to address the class imbalance. Then, they fine-tuned RoBERTa on this augmented data. **TOBB ETU** [25] used zero-shot and few-shot classification with ChatGPT exclusively for subtask 3A.

Frank [23] used CatBoost, TF.IDF, oversampling, and an ensemble. They had a file formatting issue, and thus they are not officially on the leaderboard.

5. Related Work

The detection of political bias in news articles and media has been the subject of several studies [7, 26, 27, 28, 29] due to its significance in ensuring balanced information dissemination and supporting media literacy [30].

Historically, bias was primarily understood as coverage inequality, as put forth by Stevenson et al. [31]. However, later definitions expanded to include systematic favoring of particular ideologies or candidates, as illustrated by Waldman and Devitt [32].

These broader interpretations of bias consider factors like visual favorability in news images. Based on a review of numerous studies, [33] proposed three types of media bias: gatekeeping bias, coverage bias, and statement bias. Groeling [34] influenced this classification concept of media bias, focusing on selection bias (what to cover) and presentation bias (how to cover it). Selection bias research typically involves collecting news articles or transcripts, analyzing their content, and identifying systematic biases. Meanwhile, presentation bias is often evaluated through framing, visuals, tone, and sources [35]. Multiple methods have been proposed to quantify news slant, including analyzing the language used by different political parties and mapping the distances between media sources based on their mutual followers on social media platforms like Twitter [36, 37, 29].

There have been various approaches for detecting the political bias of news articles. Kulkarni et al. [38] used an attention-based multi-view model. Baly et al. [11] used adversarial training to make sure that the model learns to predict the bias rather than the source of the news article.

Another related task is hyper-partisan news detection, e.g., [39] proposed a meta-learning approach to model the style similarities between text categories [40]. Systems using averaged word embeddings from pre-trained ELMo models succeeded for this task [41].

Efforts to predict the political ideology of news media used multimodal deep-learning Dinkov et al. [9]. Fact-checking methods that assess a document’s stance towards a claim considering the source’s credibility have also been explored [26, 28]. For assessing entire news outlets, researchers modeled tweets and Twitter users [10], information from social media, YouTube, and Wikipedia [8], and inter-media similarity based on audience overlap [12]. There have also been attempts to model bias and factuality jointly in a multitask setup [7].

The CheckThat! lab for CLEF has expanded its task offerings compared to the previous iterations, particularly concentrating on check-worthiness [42], subjectivity [43], bias (this paper), factuality [44], and authority [45]. Notably, only in this sixth edition of the CheckThat! lab a task has aimed at predicting bias at both the article level and the medium level.

Overall, various computational models and datasets have advanced the detection of political bias and factuality and contributed to media literacy efforts [46, 47, 48].

6. Conclusion and Future Work

We presented a comprehensive analysis of Task 3 from the CheckThat! lab at CLEF 2023. This lab focused on detecting the political bias in news articles and media outlets. The submissions used transformer-based models (such as RoBERTa and ChatGPT) and gradient boosting on decision trees (like CatBoost), achieving sizable improvements over the baselines.

In future work, we plan to explore more information sources, to add more languages, and to adopt a finer-grained scale.

Acknowledgments

The work of F. Alam and M. Hasanain is partially supported by NPRP 13S-0206-200281 and NPRP 14C-0916-210015 from the Qatar National Research Fund, a member of Qatar Foundation. The findings herein are solely the responsibility of the authors.

References

- [1] H. Allcott, M. Gentzkow, Social media and fake news in the 2016 election, *Journal of economic perspectives* 31 (2017) 211–236.
- [2] T. Groseclose, J. Milyo, A measure of media bias, *The quarterly journal of economics* 120 (2005) 1191–1237.
- [3] R. R. R. Gangula, S. R. Duggenpudi, R. Mamidi, Detecting political bias in news articles using headline attention, in: *Proceedings of the 2019 ACL workshop BlackboxNLP: analyzing and interpreting neural networks for NLP*, 2019, pp. 77–84.
- [4] K. Lazaridou, R. Krestel, Identifying political bias in news articles, *Bulletin of the IEEE TCDDL* 12 (2016).
- [5] R. Baly, G. Karadzhov, D. Alexandrov, J. Glass, P. Nakov, Predicting factuality of reporting and bias of news media sources, in: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, 2018, pp. 3528–3539.
- [6] R. Baly, M. Mohtarami, J. Glass, L. Màrquez, A. Moschitti, P. Nakov, Integrating stance detection and fact checking in a unified corpus, in: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, 2018, pp. 21–27.
- [7] R. Baly, G. Karadzhov, A. Saleh, J. Glass, P. Nakov, Multi-task ordinal regression for jointly predicting the trustworthiness and the leading political ideology of news media, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019, pp. 2109–2116.
- [8] R. Baly, G. Karadzhov, J. An, H. Kwak, Y. Dinkov, A. Ali, J. Glass, P. Nakov, What was written vs. who read it: News media profiling using text analysis and social media context, in: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 3364–3374.
- [9] Y. Dinkov, A. Ali, I. Koychev, P. Nakov, Predicting the leading political ideology of YouTube channels using acoustic, textual, and metadata information, in: *Proceedings of the 20th Annual Conference of the International Speech Communication Association, INTERSPEECH '19*, 2019.
- [10] P. Stefanov, K. Darwish, A. Atanasov, P. Nakov, Predicting the topical stance and political leaning of media using tweets, in: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 527–537.
- [11] R. Baly, G. Da San Martino, J. Glass, P. Nakov, We can detect your bias: Predicting the political ideology of news articles, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP '20*, 2020, pp. 4982–4991.
- [12] P. Panayotov, U. Shukla, H. T. Sencar, M. Nabeel, P. Nakov, GREENER: Graph neural networks for news media profiling, in: *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Abu Dhabi, United Arab Emirates*, 2022, pp. 7470–7480.
- [13] P. Nakov, A. Barrón-Cedeño, T. Elsayed, R. Suwaileh, L. Màrquez, W. Zaghouni, P. Atanasova, S. Kyuchukov, G. Da San Martino, Overview of the CLEF-2018 Check-That! lab on automatic identification and verification of political claims, in: *Proceedings of the Ninth International Conference of the CLEF Association: Experimental IR Meets*

Multilinguality, Multimodality, and Interaction, Lecture Notes in Computer Science, 2018, pp. 372–387.

- [14] T. Elsayed, P. Nakov, A. Barrón-Cedeño, M. Hasanain, R. Suwaileh, G. Da San Martino, P. Atanasova, Overview of the CLEF-2019 CheckThat!: Automatic identification and verification of claims, in: *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, LNCS, 2019, pp. 301–321.
- [15] A. Barrón-Cedeño, T. Elsayed, P. Nakov, G. Da San Martino, M. Hasanain, R. Suwaileh, F. Haouari, N. Babulkov, B. Hamdan, A. Nikolov, S. Shaar, Z. Ali, Overview of CheckThat! 2020 – automatic identification and verification of claims in social media, in: *Proceedings of the 11th International Conference of the CLEF Association: Experimental IR Meets Multilinguality, Multimodality, and Interaction*, CLEF '2020, 2020, pp. 215–236.
- [16] P. Nakov, G. Da San Martino, T. Elsayed, A. Barrón-Cedeño, R. Míguez, S. Shaar, F. Alam, F. Haouari, M. Hasanain, N. Babulkov, et al., The CLEF-2021 CheckThat! lab on detecting check-worthy claims, previously fact-checked claims, and fake news, in: *ECIR (2)*, 2021.
- [17] P. Nakov, A. Barrón-Cedeño, G. Da San Martino, F. Alam, J. M. Struß, T. Mandl, R. Míguez, T. Caselli, M. Kutlu, W. Zaghouni, C. Li, S. Shaar, G. K. Shahi, H. Mubarak, A. Nikolov, N. Babulkov, Y. S. Kartal, J. Beltrán, M. Wiegand, M. Siegel, J. Köhler, Overview of the CLEF-2022 CheckThat! lab on fighting the COVID-19 infodemic and fake news detection, in: A. Barrón-Cedeño, G. Da San Martino, M. Degli Esposti, F. Sebastiani, C. Macdonald, G. Pasi, A. Hanbury, M. Potthast, G. Faggioli, F. Nicola (Eds.), *Proceedings of the 13th International Conference of the CLEF Association: Information Access Evaluation meets Multilinguality, Multimodality, and Visualization*, CLEF '2022, Bologna, Italy, 2022.
- [18] P. Nakov, A. Barrón-Cedeño, G. Da San Martino, F. Alam, R. Míguez, T. Caselli, M. Kutlu, W. Zaghouni, C. Li, S. Shaar, H. Mubarak, A. Nikolov, Y. S. Kartal, J. Beltrán, Overview of the CLEF-2022 CheckThat! lab task 1 on identifying relevant claims in tweets, in: N. Faggioli, Guglielmo and Ferro, A. Hanbury, M. Potthast (Eds.), *Working Notes of CLEF 2022—Conference and Labs of the Evaluation Forum*, CLEF '2022, Bologna, Italy, 2022.
- [19] A. Barrón-Cedeño, F. Alam, T. Caselli, G. Da San Martino, T. Elsayed, A. Galassi, F. Haouari, F. Ruggeri, J. M. Struß, R. N. Nandi, G. S. Cheema, D. Azizov, P. Nakov, The CLEF-2023 CheckThat! Lab: Checkworthiness, subjectivity, political bias, factuality, and authority, in: J. Kamps, L. Goeuriot, F. Crestani, M. Maistro, H. Joho, B. Davis, C. Gurrin, U. Kruschwitz, A. Caputo (Eds.), *Advances in Information Retrieval*, Springer Nature Switzerland, Cham, 2023, pp. 506–517.
- [20] A. Barrón-Cedeño, T. Elsayed, P. Nakov, G. Da San Martino, M. Hasanain, R. Suwaileh, F. Haouari, Checkthat! at clef 2020: Enabling the automatic identification and verification of claims in social media, in: J. M. Jose, E. Yilmaz, J. Magalhães, P. Castells, N. Ferro, M. J. Silva, F. Martins (Eds.), *Advances in Information Retrieval*, Cham, 2020, pp. 499–507.
- [21] A. Barrón-Cedeño, F. Alam, A. Galassi, G. Da San Martino, P. Nakov, , T. Elsayed, D. Azizov, T. Caselli, G. Cheema, F. Haouari, M. Hasanain, M. Kutlu, C. Li, F. Ruggeri, J. M. Struß, W. Zaghouni, Overview of the CLEF–2023 CheckThat! Lab checkworthiness, subjectivity, political bias, factuality, and authority of news articles and their source, in: A. Arampatzis, E. Kanoulas, T. Tsikrika, S. Vrochidis, A. Giachanou, D. Li, M. Aliannejadi, M. Vlachos, G. Faggioli, N. Ferro (Eds.), *Experimental IR Meets Multilinguality, Multimodality, and Interaction*. *Proceedings of the Fourteenth International Conference of the CLEF Association*

(CLEF 2023), 2023.

- [22] G. Da San Martino, F. Alam, M. Hasanain, R. N. Nandi, D. Azizov, P. Nakov, Overview of the CLEF-2023 CheckThat! lab task 3 on political bias of news articles and news media, in: Working Notes of CLEF 2023–Conference and Labs of the Evaluation Forum, CLEF '2023, Thessaloniki, Greece, 2023.
- [23] D. Azizov, S. Liang, P. Nakov, Frank at checkthat! 2023: Detecting the political bias of news articles and news media, in: Working Notes of CLEF 2023–Conference and Labs of the Evaluation Forum, CLEF '2023, Thessaloniki, Greece, 2023.
- [24] S. Tran, P. Rodrigues, B. Strauss, E. Williams, Accenture at CheckThat! 2023: Learning to detect Political Bias of News Articles and Sources, in: [49], 2023.
- [25] M. D. Türkmen, G. Coşgun, M. Kutlu, Tobb etu at checkthat! 2023: Utilizing chatgpt to detect subjective statements and political bias, in: [49], 2023.
- [26] R. Baly, G. Karadzhov, D. Alexandrov, J. Glass, P. Nakov, Predicting factuality of reporting and bias of news media sources, in: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, 2018, pp. 3528–3539.
- [27] R. Baly, G. Da San Martino, J. Glass, P. Nakov, We can detect your bias: Predicting the political ideology of news articles, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020, pp. 4982–4991.
- [28] R. Baly, M. Mohtarami, J. Glass, L. Màrquez, A. Moschitti, P. Nakov, Integrating stance detection and fact checking in a unified corpus, in: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers), New Orleans, Louisiana, 2018, pp. 21–27.
- [29] P. Nakov, H. T. Sencar, J. An, H. Kwak, A survey on predicting the factuality and the bias of news media, arXiv preprint arXiv:2103.12506 (2021).
- [30] M. Gentzkow, J. M. Shapiro, What drives media slant? Evidence from US daily newspapers, *Econometrica* 78 (2010) 35–71.
- [31] R. L. Stevenson, R. A. Eisinger, B. M. Feinberg, A. B. Kotok, Untwisting the news twist: A replication of efron’s study, *Journalism Quarterly* 50 (1973) 211–219.
- [32] P. Waldman, J. Devitt, Newspaper photographs and the 1996 presidential election: The question of bias, *Journalism & Mass Communication Quarterly* 75 (1998) 302–311.
- [33] D. D’Alessio, M. Allen, Media bias in presidential elections: A meta-analysis, *Journal of communication* 50 (2000) 133–156.
- [34] T. Groeling, Media bias by the numbers: Challenges and opportunities in the empirical study of partisan news, *Annual Review of Political Science* 16 (2013) 129–151.
- [35] E. C. Tandoc Jr, Journalism is twerking? How web analytics is changing the process of gatekeeping, *New media & society* 16 (2014) 559–575.
- [36] K. Darwish, P. Stefanov, M. Aupetit, P. Nakov, Unsupervised user stance detection on Twitter, in: Proceedings of the International AAAI Conference on Web and Social Media, volume 14, 2020, pp. 141–152.
- [37] P. Stefanov, K. Darwish, A. Atanasov, P. Nakov, Predicting the topical stance and political leaning of media using tweets, in: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, 2020, pp. 527–537.
- [38] V. Kulkarni, J. Ye, S. Skiena, W. Y. Wang, Multi-view models for political ideology detection of news articles, in: Proceedings of the 2018 Conference on Empirical Methods in Natural

- Language Processing, 2018, pp. 3518–3527.
- [39] J. Kiesel, M. Mestre, R. Shukla, E. Vincent, P. Adineh, D. Corney, B. Stein, M. Potthast, Semeval-2019 task 4: Hyperpartisan news detection, in: Proceedings of the 13th International Workshop on Semantic Evaluation, 2019, pp. 829–839.
 - [40] M. Potthast, M. Hagen, T. Gollub, M. Tippmann, J. Kiesel, P. Rosso, E. Stamatatos, B. Stein, Overview of the 5th international competition on plagiarism detection, in: CLEF Conference on Multilingual and Multimodal Information Access Evaluation, CELCT, 2013, pp. 301–331.
 - [41] Y. Jiang, J. Petrak, X. Song, K. Bontcheva, D. Maynard, Team berthava von suttner at SemEval-2019 task 4: Hyperpartisan news detection using ELMo sentence representation convolutional network, in: Proceedings of the 13th International Workshop on Semantic Evaluation, Minneapolis, Minnesota, USA, 2019, pp. 840–844.
 - [42] F. Alam, A. Barrón-Cedeño, G. S. Cheema, S. Hakimov, M. Hasanain, C. Li, R. Míguez, H. Mubarak, G. K. Shahi, W. Zaghouni, P. Nakov, Overview of the CLEF-2023 CheckThat! lab task 1 on check-worthiness in multimodal and multigenre content, in: Working Notes of CLEF 2023–Conference and Labs of the Evaluation Forum, CLEF '2023, Thessaloniki, Greece, 2023.
 - [43] A. Galassi, F. Ruggeri, A. B.-C. no, F. Alam, T. Caselli, M. Kutlu, J. M. Struss, F. Antici, M. Hasanain, J. Köhler, K. Korre, F. Leistra, A. Muti, M. Siegel, M. D. Turkmen, M. Wiegand, W. Zaghouni, Overview of the CLEF-2023 CheckThat! lab task 2 on subjectivity in news articles, in: Working Notes of CLEF 2023–Conference and Labs of the Evaluation Forum, CLEF '2023, Thessaloniki, Greece, 2023.
 - [44] P. Nakov, F. Alam, G. Da San Martino, M. Hasanain, R. N. Nandi, D. Azizov, P. Panayotov, Overview of the CLEF-2023 CheckThat! lab task 4 on factuality of reporting of news media, in: Working Notes of CLEF 2023–Conference and Labs of the Evaluation Forum, CLEF '2023, Thessaloniki, Greece, 2023.
 - [45] F. Haouari, Z. Sheikh Ali, T. Elsayed, Overview of the CLEF-2023 CheckThat! lab task 5 on authority finding in Twitter, in: Working Notes of CLEF 2023–Conference and Labs of the Evaluation Forum, CLEF '2023, Thessaloniki, Greece, 2023.
 - [46] P. Nakov, D. Corney, M. Hasanain, F. Alam, T. Elsayed, A. Barrón-Cedeño, P. Papotti, S. Shaar, G. D. S. Martino, Automated fact-checking for assisting human fact-checkers, in: Proceedings of the 30th International Joint Conference on Artificial Intelligence, 2021.
 - [47] F. Alam, S. Shaar, F. Dalvi, H. Sajjad, A. Nikolov, H. Mubarak, G. D. S. Martino, A. Abdelali, N. Durrani, K. Darwish, A. Al-Homaid, W. Zaghouni, T. Caselli, G. Danoe, F. Stolk, B. Bruntink, P. Nakov, Fighting the COVID-19 infodemic: Modeling the perspective of journalists, fact-checkers, social media platforms, policy makers, and the society, in: Findings of EMNLP 2021, 2021, pp. 611–649.
 - [48] P. Nakov, S. Rosenthal, Z. Kozareva, V. Stoyanov, A. Ritter, T. Wilson, SemEval-2013 task 2: Sentiment analysis in Twitter, in: Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013), 2013, pp. 312–320.
 - [49] M. Aliannejadi, G. Faggioli, N. Ferro, Vlachos, Michalis (Eds.), Working Notes of CLEF 2023 - Conference and Labs of the Evaluation Forum, CLEF 2023, Thessaloniki, Greece, 2023.