



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE
DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

The CLEF-2022 CheckThat! Lab on Fighting the COVID-19 Infodemic and Fake News Detection

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

The CLEF-2022 CheckThat! Lab on Fighting the COVID-19 Infodemic and Fake News Detection / Nakov, Preslav; Barron Cedeno, Alberto; Da San Martino, Giovanni; Alam, Firoj; Struß, Julia Maria; Mandl, Thomas; Míguez, Rubén; Caselli, Tommaso; Kutlu, Mucahid; Zaghouani, Wajdi; Li, Chengkai; Shaar, Shaden; Shahi, Gautam Kishore; Mubarak, Hamdy; Nikolov, Alex; Babulkov, Nikolay; Kartal, Yavuz Selim; Beltrán, Javier. - ELETTRONICO. - 13186:(2022), pp. 416-428. (Intervento presentato al convegno European Conference on Information Retrieval tenutosi a Stavanger, Norway nel 10/04/2022 - 14/04/2022) [10.1007/978-3-030-99739-7_52]
This version is available at: <https://hdl.handle.net/11585/903214> since: 2022-11-16

Published:

DOI: http://doi.org/10.1007/978-3-030-99739-7_52

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Nakov, Barrón-Cedeño, Da San Martino, Alam, Struß, Mandl, Míguez, Caselli, Kutlu, Zaghouani, Li, Shaar, Shahi, Mubarak, Nikolov, Babulkov, Kartal, Beltrán (2022). The CLEF-2022 CheckThat! Lab on Fighting the COVID-19 Infodemic and Fake News Detection. In: , *et al.* Advances in Information Retrieval. ECIR 2022. Lecture Notes in Computer Science, vol 13186. Springer, Cham. **The final published version is available online at:** https://doi.org/10.1007/978-3-030-99739-7_52

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

The CLEF-2022 CheckThat! Lab on Fighting the COVID-19 Infodemic and Fake News Detection

Preslav Nakov¹, Alberto Barrón-Cedeño², Giovanni da San Martino³,
Firoj Alam¹, Julia Maria Struß⁴, Thomas Mandl⁵, Rubén Míguez⁶,
Tommaso Caselli⁷, Mucahid Kutlu⁸, Wajdi Zaghouani⁹, Chengkai Li¹⁰,
Shaden Shaar¹¹, Gautam Kishore Shahi¹², Hamdy Mubarak¹, Alex Nikolov¹³,
Nikolay Babulkov¹⁴, Yavuz Selim Kartal¹⁵, Javier Beltrán⁶

¹Qatar Computing Research Institute, HBKU, ²DIT, Università di Bologna,
³University of Padova, ⁴University of Applied Sciences Potsdam, ⁵University
of Hildesheim, ⁶Newtral Media Audiovisual, ⁷University of Groningen,
⁸TOBB University of Economics and Technology, ⁹Hamad Bin Khalifa
University, ¹⁰University of Texas at Arlington, ¹¹Cornell University,
¹²University of Duisburg-Essen, ¹³CheckStep, ¹⁴Sofia University,
¹⁵GESIS – Leibniz Institute for the Social Sciences

{pnakov, fialam, hmubarak}@hbku.edu.qa, a.barron@unibo.it,
dasan@math.unipd.it, struss@fh-potsdam.de, mandl@uni-hildesheim.de,
{ruben.miguez, javier.beltran}@newtral.es, m.kutlu@etu.edu.tr,
wzaghouani@hbku.edu.qa, ss2753@cornell.edu, gautam.shahi@uni-due.de,
alex.nikolov@checkstep.com, yavuzselim.kartal@gesis.org

Abstract. The fifth edition of the CheckThat! Lab is held as part of the 2022 Conference and Labs of the Evaluation Forum (CLEF). The lab evaluates technology supporting various factuality tasks in seven languages: Arabic, Bulgarian, Dutch, English, German, Spanish, and Turkish. Task 1 focuses on disinformation related to the ongoing COVID-19 infodemic and politics, and asks to predict whether a tweet is worth fact-checking, contains a verifiable factual claim, is harmful to the society, or is of interest to policy makers and why. Task 2 asks to retrieve claims that have been previously fact-checked and that could be useful to verify the claim in a tweet. Task 3 is to predict the veracity of a news article. Tasks 1 and 3 are classification problems, while Task 2 is a ranking one.

Keywords: Fact-Checking, Disinformation, Misinformation, Check-Worthiness, Verified Claim Retrieval, Fake News, Factuality, COVID-19.

1 Introduction

The mission of the CheckThat! lab is to foster the development of technology to assist in the process of fact-checking news articles, political debates, and social media posts. Four editions of the lab have been held previously, targeting various natural language processing and information retrieval tasks related to factuality.

The 2018 edition of the lab focused on check-worthiness and fact-checking of claims in political debates [54]. The 2019 edition covered the various modules necessary to verify a claim: from check-worthiness, to ranking and classification of evidence in the form of Web pages, to actual fact-checking of claims against specific text snippets [24, 25]. The 2020 edition featured three main tasks: detecting previously fact-checked claims, evidence retrieval, and actual fact-checking of claims [9, 11]. Similarly, the 2021 edition focused on detecting check-worthy claims, previously fact-checked claims, and fake news [56, 57]. Whereas the first editions focused mostly on political debates and speeches, and eventually tweets, the 2021 edition added the verification of news articles. Notably, all editions covered one of the most important initial stages in the fact-checking process: the identification of check-worthy claims—in debates, speeches, press conferences, and tweets. Finally, over the years, **CheckThat!** has witnessed an expansion in terms of language coverage, going from two (Arabic and English) to seven languages now (Arabic, Bulgarian, Dutch, English, German, Spanish, and Turkish).

The 2022 edition of the lab features three tasks to foster the technology on three timely problems in multiple languages.¹ Task 1 asks to detect relevant tweets: check-worthy, verifiable, harmful, and attention-worthy. Task 2 aims at detecting previously fact-checking claims. Task 3 focuses on checking the factuality of news articles. Automated systems to detect and to verify such multifaceted aspects can be very useful as supportive technology for investigative journalism, as they could provide help and guidance, thus saving time [3, 27, 37, 39, 55, 81]. For example, a system could automatically identify check-worthy claims, could make sure they have not been fact-checked before by a reputable fact-checking organization, and can then present them to a journalist for further analysis in a ranked list. Similarly, a system can identify harmful and attention-worthy social media content to support different stakeholders in their day-to-day decision-making process.

2 Description of the Tasks

The lab is organized around three tasks, each of which in turn has several subtasks. Figure 1 shows the full **CheckThat!** verification pipeline, with the three tasks we target this year highlighted.

2.1 Task 1: Identifying Relevant Claims in Tweets

Task 1 has four subtasks, three binary and one multi-class; Table 1 shows the class labels for each task. More detail about the original dataset on which this task is based can be found in [2, 3].

Subtask 1A: Check-Worthiness Estimation Given a tweet, predict whether it is worth fact-checking by professional fact-checkers.

Subtask 1B: Verifiable Factual Claims Detection. Given a tweet, predict whether it contains a verifiable factual claim.

¹ <http://sites.google.com/view/clef2022-checkthat/>




Fig. 1. The full verification pipeline. The lab covers three tasks from this pipeline: 1. identifying relevant claims in tweets, 2. verified claim retrieval, and 3. fake news detection. The languages covered are Arabic (ar), Bulgarian (bg), Dutch (nl), English (en), German (de), Spanish (es), and Turkish (tr).

Subtask 1A	Subtask 1C	Subtask 1D	
1. No	1. No	1. No	6. Yes, contains advice
2. Yes	2. Yes	2. Yes, asks question	7. Yes, discusses action taken
Subtask 1B		3. Yes, blame authorities	8. Yes, discusses cure
1. No		4. Yes, calls for action	9. Yes, other
2. Yes		5. Yes, Harmful	

Table 1. Overview of the classes for Subtasks 1A, 1B, 1C, and 1D.

Subtask 1C: Harmful Tweet Detection. Given a tweet, predict whether it is harmful to the society.

Subtask 1D: Attention-Worthy Tweet Detection. Given a tweet, predict whether it should get the attention of policy makers and why.

2.2 Task 2: Detecting Previously Fact-Checked Claims

Given a check-worthy claim, and a set of previously-checked claims, determine whether the claim has been previously fact-checked with respect to a collection of fact-checked claims [67, 69].

Subtask 2A: Detecting previously fact-checked claims from tweets. Given a tweet, detect whether the claim the tweet makes was previously fact-checked with respect to a collection of previously fact-checked claims. This is a ranking task, where the systems are asked to produce a list of top- n candidates.

Subtask 2B: Detecting previously fact-checked claims in political debates/speeches. Within the context of a political debate or a speech, detect whether a claim has been previously fact-checked with respect to a collection of previously fact-checked claims. This is a ranking task, where systems are asked to produce a list of top- n candidates. It is offered in English only.

Task	ar	bl	en	nl	es	tr	Task	ar	en	Task	en	de
1A	4.1k	2.6k	3.0k	1.9k	15.0k	3.6k	2A: Input-claim	908	1.6k	2B: Input-claim	752	
1B	6.2k	4.0k	4.7k	4.0k		3.8k	2A: Input-VC	1.0k	1.6k	2B: Input-VC	869	
1C	6.1k	4.0k	4.7k	4.0k		3.8k	2A: VC	30.3k	13.8k	2B: VC	20.7k	
1D	6.1k	4.0k	4.7k	4.0k		3.8k				3	1.8k	650

*The process of crawling and annotating the data is not finished, and the final number will vary.
Table 2. Data for all tasks. *VC*: Verified claim, *Input-VC*: Input-verified claim pair.

2.3 Task 3: Fake News Detection

This task targets news articles. Given the text and the title of an article, determine whether the main claim made in the article is true, partially true, false, or other (e.g., articles in dispute and unproven articles) [75, 76]. This task is offered as a monolingual task in English and an English-German cross-lingual task.

3 Data

Table 2 summarizes the data available for each task and for each language.

Task 1: Identifying Relevant Claims in Tweets. We have more than 34K annotations about several topics, including COVID-19 and politics, which cover all subtasks 1A-1D [10, 57].

Task 2: Detecting Previously Fact-Checked Claims. For Subtask 2A, we have 1,400 annotated examples ranging from 2016 till 2021. The test set comes from Snopes. For more details, refer to [71]. For Subtask 2B, we have 800 claims available for training [70].

Task 3: Fake News Detection We have 1,254 annotated examples in English on various topics including COVID-19, climate change, and politics from the 2021 edition of the lab [77]. We provide a new test set for English and a similarly annotated test data set in German.

4 Evaluation Measures

For the classification tasks 1A and 1C, we use the F_1 measure with respect to the positive class, for Task 1B, we use accuracy, and for Task 1D, we use weighted- F_1 . For the ranking problems in Tasks 2A and 2B, the official evaluation measure is *Mean Average Precision* (MAP), as in the two previous editions of these tasks. We also report reciprocal rank, and $P@k$ for $k \in \{1, 3, 5, 10, 20, 30\}$. Finally, for Task 3, we use macro-average F_1 -measure as the official evaluation measure.

5 Previously in the CheckThat! Lab

Four editions of **CheckThat!** have been held so far, and some of the tasks in the 2022 edition are closely related to tasks from previous editions. Thus, considering the most successful approaches applied in the past is a good starting point to address the 2022 tasks. Below, we briefly discuss the tasks from previous years.

5.1 CheckThat! 2021

Task 1₂₀₂₁. Given a topic and a set of potentially related tweets (or political debates/speeches), rank the tweets (or the sentences in the political debate/speech) according to their check-worthiness for the topic. BERT, AraBERT, and RoBERTa were by far the most popular large-scale pre-trained language models for the task [72, 83]. Other approaches used WordNet [85] and LIWC [66].

Task 2₂₀₂₁. Given a tweet, a political debate, or a speech, detect whether the claim it makes was previously fact-checked with respect to a collection of fact-checked claims. The most successful approaches were based on AraBERT, RoBERTa, and Sentence-BERT [18, 47, 65].

Task 3₂₀₂₁. Given the text and the title of a news article, determine whether the main claim it makes is true, partially true, false, or other. Also, identify the domain of the article: health, crime, climate, elections, or education. The most successful pre-trained language model was RoBERTa [7, 19, 44]. Ensembles were also popular, with components using BERT [44] and LSTMs [19, 44].

5.2 CheckThat! 2020

Task 1₂₀₂₀. Given a topic and a stream of potentially-related tweets, rank the tweets by check-worthiness for the topic [35, 73]. The most successful runs adopted state-of-the-art transformers. The top-ranked teams for the English version of this task used BERT [17] and RoBERTa [60, 82]. For the Arabic version, the top systems used AraBERT [42, 82] and a multilingual BERT [34].

Task 2₂₀₂₀. Given a check-worthy claim and a collection of previously verified claims, rank these verified claims, so that those that verify the input claim (or a sub-claim in it) are ranked on top [73]. The most effective approaches fine-tuned BERT or RoBERTa [13].

Task 3₂₀₂₀. Given a check-worthy claim in a tweet on a specific topic and a set of text snippets extracted from potentially relevant Web pages, return a ranked list of evidence snippets for the claim.

Task 4₂₀₂₀. Given a check-worthy claim on a specific topic and a set of potentially-relevant Web pages, predict the veracity of the claim [35]. The top model used a scoring function that computes the degree of concordance and negation between a claim and all input text snippets for that claim [80].

Task 5₂₀₂₀. Given a debate segmented into sentences, together with speaker information, prioritize sentences for fact-checking [73]. Only one out of eight runs outperformed a strong bi-LSTM baseline [46].

5.3 CheckThat! 2019

Task 1₂₀₁₉. Given a political debate, an interview, or a speech, rank its sentences by the priority with which they should be fact-checked [5]. The most successful approaches used neural networks for the individual classification of the instances, e.g., based on domain-specific word embeddings, syntactic dependencies, and LSTMs [33].

Task 2₂₀₁₉. Given a claim and a set of potentially relevant Web pages, identify which of these pages (and passages thereof) are useful for a human to fact-check the claim. Finally, determine the factuality of the claim [36]. The best approach used textual entailment and external data [28].

5.4 CheckThat! 2018

Task 1₂₀₁₈. It was identical to Task 1₂₀₁₉ [4]. The best approaches used *pseudo-speeches* as a concatenation of all interventions by a debater [87], and represented the entries with embeddings, POS tags, and syntactic dependencies [32].

Task 2₂₀₁₈. Given a check-worthy claim in the form of a (transcribed) sentence, determine whether the claim is true, half-true, or false [12]. The best approach grabbed information from the Web and fed the claim with the most similar Web text into a CNN [32].

6 Related Work

There has been a lot of research on checking the factuality of a claim, of a news article, or of an information source [6, 8, 41, 45, 51, 59, 64, 68, 86]. Special attention has been paid to disinformation and misinformation in social media [30, 43, 49, 74, 78, 84], more recently with focus on fighting the COVID-19 infodemic [2, 3, 52, 53]. Check-worthiness estimation is still an understudied problem, especially in social media [27, 37–40, 81], and fake news detection for news articles is mostly approached as a binary classification problem [59, 61].

CheckThat! is related to several tasks at SemEval: on determining rumour veracity [22, 29], on stance detection [50], on fact-checking in community question answering forums [48], on propaganda detection [21, 23], and on semantic textual similarity [1, 58]. It is also related to the FEVER task [79] on fact extraction and verification, to the Fake News Challenge [31, 63], to the FakeNews task at MediaEval [62], as well as to the NLP4IF tasks on propaganda detection [20] and on fighting the COVID-19 infodemic in social media [68].

7 Conclusion

We presented the 2022 edition of the **CheckThat!** lab, which features tasks that span the full fact-checking pipeline: from spotting check-worthy claims to checking whether an input claim has been fact-checked before. We further have a fake news detection task. Last but not least, in line with one of the main missions of CLEF, we promote multi-linguality by offering tasks in seven languages: Arabic, Bulgarian, Dutch, English, German, Spanish, and Turkish.

Acknowledgments

This research is part of the Tanbih mega-project, developed at the Qatar Computing Research Institute, HBKU, which aims to limit the impact of “fake news”, propaganda, and media bias, thus promoting media literacy and critical thinking.

References

1. Agirre, E., Banea, C., Cer, D., Diab, M., Gonzalez-Agirre, A., Mihalcea, R., Rigau, G., Wiebe, J.: SemEval-2016 task 1: Semantic textual similarity, monolingual and cross-lingual evaluation. In: Proceedings of the 10th International Workshop on Semantic Evaluation. pp. 497–511. SemEval '16 (2016)
2. Alam, F., Dalvi, F., Shaar, S., Durrani, N., Mubarak, H., Nikolov, A., Da San Martino, G., Abdelali, A., Sajjad, H., Darwish, K., Nakov, P.: Fighting the COVID-19 infodemic in social media: A holistic perspective and a call to arms. In: Proceedings of the International AAAI Conference on Web and Social Media. pp. 913–922. ICWSM '21 (2021)
3. Alam, F., Shaar, S., Dalvi, F., Sajjad, H., Nikolov, A., Mubarak, H., Martino, G.D.S., Abdelali, A., Durrani, N., Darwish, K., Al-Homaid, A., Zaghoulani, W., Caselli, T., Danoe, G., Stolk, F., Bruntink, B., Nakov, P.: 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. pp. 611–649 (2021)
4. Atanasova, P., Marquez, L., Barrón-Cedeño, A., Elsayed, T., Suwaileh, R., Zaghoulani, W., Kyuchukov, S., Da San Martino, G., Nakov, P.: Overview of the CLEF-2018 CheckThat! lab on automatic identification and verification of political claims. Task 1: Check-worthiness. In: Cappellato et al. [16]
5. Atanasova, P., Nakov, P., Karadzhov, G., Mohtarami, M., Da San Martino, G.: Overview of the CLEF-2019 CheckThat! lab on automatic identification and verification of claims. Task 1: Check-worthiness. In: Cappellato et al. [15]
6. Ba, M.L., Berti-Equille, L., Shah, K., Hammady, H.M.: VERA: A platform for veracity estimation over web data. In: Proceedings of the 25th International Conference on World Wide Web. pp. 159–162. WWW '16 (2016)
7. Balouchzahi, F., Shashirekha, H., Sidorov, G.: MUCIC at CheckThat! 2021: FaDo-fake news detection and domain identification using transformers ensembling. In: Faggioli et al. [26], pp. 455–464
8. 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. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 3364–3374. ACL '20 (2020)
9. Barrón-Cedeño, A., Elsayed, T., Nakov, P., Da San Martino, G., Hasanain, M., Suwaileh, R., Haouari, F.: CheckThat! at CLEF 2020: Enabling the automatic identification and verification of claims in social media. In: Advances in Information Retrieval. pp. 499–507. Springer International Publishing, Cham (2020)
10. Barrón-Cedeño, A., Elsayed, T., Nakov, P., Da San Martino, G., Hasanain, M., Suwaileh, R., Haouari, F., Babulkov, N., Hamdan, B., Nikolov, A., Shaar, S., Ali, Z.: 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. pp. 215–236. CLEF '2020 (2020)
11. Barrón-Cedeño, A., Elsayed, T., Nakov, P., Da San Martino, G., Hasanain, M., Suwaileh, R., Haouari, F., Babulkov, N., Hamdan, B., Nikolov, A., Shaar, S., Sheikh Ali, Z.: Overview of CheckThat! 2020: Automatic identification and verification of claims in social media. pp. 215–236. LNCS (12260), Springer (2020)
12. Barrón-Cedeño, A., Elsayed, T., Suwaileh, R., Marquez, L., Atanasova, P., Zaghoulani, W., Kyuchukov, S., Da San Martino, G., Nakov, P.: Overview of the

- CLEF-2018 CheckThat! lab on automatic identification and verification of political claims. Task 2: Factuality. In: Cappellato et al. [16]
13. Bouziane, M., Perrin, H., Cluzeau, A., Mardas, J., Sadeq, A.: Buster.AI at CheckThat! 2020: Insights and recommendations to improve fact-checking. In: Cappellato et al. [14]
 14. Cappellato, L., Eickhoff, C., Ferro, N., Névéal, A. (eds.): CLEF 2020 Working Notes. CEUR Workshop Proceedings (2020)
 15. Cappellato, L., Ferro, N., Losada, D., Müller, H. (eds.): Working Notes of CLEF 2019 Conference and Labs of the Evaluation Forum. CEUR Workshop Proceedings (2019)
 16. Cappellato, L., Ferro, N., Nie, J.Y., Soulier, L. (eds.): Working Notes of CLEF 2018–Conference and Labs of the Evaluation Forum. CEUR Workshop Proceedings (2018)
 17. Cheema, G.S., Hakimov, S., Ewerth, R.: Check_square at CheckThat! 2020: Claim detection in social media via fusion of transformer and syntactic features. In: Cappellato et al. [14]
 18. Chernyavskiy, A., Ilvovsky, D., Nakov, P.: Aschern at CLEF CheckThat! 2021: Lambda-calculus of fact-checked claims. In: Faggioli et al. [26]
 19. Cusmulic, C.G., Amarandei, M.A., Pelin, I., Cociorva, V.I., Iftene, A.: UAICS at CheckThat! 2021: Fake news detection. In: Faggioli et al. [26]
 20. Da San Martino, G., Barrón-Cedeño, A., Nakov, P.: Findings of the NLP4IF-2019 shared task on fine-grained propaganda detection. In: Proceedings of the Second Workshop on Natural Language Processing for Internet Freedom: Censorship, Disinformation, and Propaganda. pp. 162–170. NLP4IF '19 (2019)
 21. Da San Martino, G., Barrón-Cedeño, A., Wachsmuth, H., Petrov, R., Nakov, P.: SemEval-2020 task 11: Detection of propaganda techniques in news articles. In: Proceedings of the 14th Workshop on Semantic Evaluation. pp. 1377–1414. SemEval '20 (2020)
 22. Derczynski, L., Bontcheva, K., Liakata, M., Procter, R., Wong Sak Hoi, G., Zubiaga, A.: SemEval-2017 task 8: RumourEval: Determining rumour veracity and support for rumours. In: Proceedings of the 11th International Workshop on Semantic Evaluation. pp. 69–76. SemEval '17 (2017)
 23. Dimitrov, D., Ali, B.B., Shaar, S., Alam, F., Silvestri, F., Firooz, H., Nakov, P., Martino, G.D.S.: SemEval-2021 task 6: Detection of persuasion techniques in texts and images. In: Proceedings of the International Workshop on Semantic Evaluation. pp. 70–98. SemEval '21 (2021)
 24. Elsayed, T., Nakov, P., Barrón-Cedeño, A., Hasanain, M., Suwaileh, R., Da San Martino, G., Atanasova, P.: CheckThat! at CLEF 2019: Automatic identification and verification of claims. In: Azzopardi, L., Stein, B., Fuhr, N., Mayr, P., Hauff, C., Hiemstra, D. (eds.) *Advances in Information Retrieval*. pp. 309–315. Springer International Publishing, Cham (2019)
 25. Elsayed, T., Nakov, P., Barrón-Cedeño, A., Hasanain, M., Suwaileh, R., Da San Martino, G., Atanasova, P.: Overview of the CLEF-2019 CheckThat!: Automatic identification and verification of claims. In: *Experimental IR Meets Multilinguality, Multimodality, and Interaction*. pp. 301–321. LNCS (2019)
 26. Faggioli, G., Ferro, N., Joly, A., Maistro, M., Piroi, F. (eds.): CLEF 2021 Working Notes. Working Notes of CLEF 2021–Conference and Labs of the Evaluation Forum (2021)
 27. Gencheva, P., Nakov, P., Márquez, L., Barrón-Cedeño, A., Koychev, I.: A context-aware approach for detecting worth-checking claims in political debates. In: Pro-

- ceedings of the International Conference Recent Advances in Natural Language Processing. pp. 267–276. RANLP’17 (2017)
28. Ghanem, B., Glavaš, G., Giachanou, A., Ponzetto, S., Rosso, P., Rangel, F.: UPV-UMA at CheckThat! lab: Verifying Arabic claims using cross lingual approach. In: Cappellato et al. [15]
 29. Gorrell, G., Kochkina, E., Liakata, M., Aker, A., Zubiaga, A., Bontcheva, K., Derczynski, L.: SemEval-2019 task 7: RumourEval, determining rumour veracity and support for rumours. In: Proceedings of the 13th International Workshop on Semantic Evaluation. pp. 845–854. SemEval ’19 (2019)
 30. Gupta, A., Kumaraguru, P., Castillo, C., Meier, P.: TweetCred: Real-time credibility assessment of content on Twitter. In: Proceedings of the 6th International Social Informatics Conference. pp. 228–243. SocInfo ’14 (2014)
 31. Hanselowski, A., PVS, A., Schiller, B., Caspelherr, F., Chaudhuri, D., Meyer, C.M., Gurevych, I.: A retrospective analysis of the fake news challenge stance-detection task. In: Proceedings of the 27th International Conference on Computational Linguistics. pp. 1859–1874. COLING ’18 (2018)
 32. Hansen, C., Hansen, C., Simonsen, J., Lioma, C.: The Copenhagen team participation in the check-worthiness task of the competition of automatic identification and verification of claims in political debates of the CLEF-2018 fact checking lab. In: Cappellato et al. [16]
 33. Hansen, C., Hansen, C., Simonsen, J., Lioma, C.: Neural weakly supervised fact check-worthiness detection with contrastive sampling-based ranking loss. In: Cappellato et al. [15]
 34. Hasanain, M., Elsayed, T.: bigIR at CheckThat! 2020: Multilingual BERT for ranking Arabic tweets by check-worthiness. In: Cappellato et al. [14]
 35. Hasanain, M., Haouari, F., Suwaileh, R., Ali, Z., Hamdan, B., Elsayed, T., Barrón-Cedeño, A., Da San Martino, G., Nakov, P.: Overview of CheckThat! 2020 Arabic: Automatic identification and verification of claims in social media. In: Cappellato et al. [14]
 36. Hasanain, M., Suwaileh, R., Elsayed, T., Barrón-Cedeño, A., Nakov, P.: Overview of the CLEF-2019 CheckThat! lab on automatic identification and verification of claims. Task 2: Evidence and factuality. In: Cappellato et al. [15]
 37. Hassan, N., Li, C., Tremayne, M.: Detecting check-worthy factual claims in presidential debates. In: Proceedings of the 24th ACM International on Conference on Information and Knowledge Management. pp. 1835–1838. CIKM ’15 (2015)
 38. Hassan, N., Tremayne, M., Arslan, F., Li, C.: Comparing automated factual claim detection against judgments of journalism organizations. In: Computation+Journalism Symposium. pp. 1–5 (2016)
 39. Hassan, N., Zhang, G., Arslan, F., Caraballo, J., Jimenez, D., Gawsane, S., Hasan, S., Joseph, M., Kulkarni, A., Nayak, A.K., et al.: ClaimBuster: The first-ever end-to-end fact-checking system. Proceedings of the VLDB Endowment **10**(12), 1945–1948 (2017)
 40. Jaradat, I., Gencheva, P., Barrón-Cedeño, A., Márquez, L., Nakov, P.: ClaimRank: Detecting check-worthy claims in Arabic and English. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations. pp. 26–30. NAACL-HLT ’18 (2018)
 41. Karadzhev, G., Nakov, P., Márquez, L., Barrón-Cedeño, A., Koychev, I.: Fully automated fact checking using external sources. In: Proceedings of the International Conference Recent Advances in Natural Language Processing. pp. 344–353. RANLP’ 17 (2017)

42. Kartal, Y.S., Kutlu, M.: TOBB ETU at CheckThat! 2020: Prioritizing English and Arabic claims based on check-worthiness. In: Cappellato et al. [14]
43. Kazemi, A., Garimella, K., Gaffney, D., Hale, S.: Claim matching beyond English to scale global fact-checking. In: Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing. pp. 4504–4517. ACL-IJCNLP '21 (2021)
44. Kovachevich, N.: BERT fine-tuning approach to CLEF CheckThat! fake news detection. In: Faggioli et al. [26]
45. Ma, J., Gao, W., Mitra, P., Kwon, S., Jansen, B.J., Wong, K.F., Cha, M.: Detecting rumors from microblogs with recurrent neural networks. In: Proceedings of the International Joint Conference on Artificial Intelligence. pp. 3818–3824. IJCAI '16 (2016)
46. Martinez-Rico, J., Araujo, L., Martinez-Romo, J.: NLP&IR@UNED at CheckThat! 2020: A preliminary approach for check-worthiness and claim retrieval tasks using neural networks and graphs. In: Cappellato et al. [14]
47. Mihaylova, S., Borisova, I., Chemishanov, D., Hadzhitsanev, P., Hardalov, M., Nakov, P.: DIPS at CheckThat! 2021: Verified claim retrieval. In: Faggioli et al. [26]
48. Mihaylova, T., Karadzhov, G., Atanasova, P., Baly, R., Mohtarami, M., Nakov, P.: SemEval-2019 task 8: Fact checking in community question answering forums. In: Proceedings of the 13th International Workshop on Semantic Evaluation. pp. 860–869. SemEval '19 (2019)
49. Mitra, T., Gilbert, E.: CREDBANK: A large-scale social media corpus with associated credibility annotations. In: Proceedings of the Ninth International AAAI Conference on Web and Social Media. pp. 258–267. ICWSM '15 (2015)
50. Mohammad, S., Kiritchenko, S., Sobhani, P., Zhu, X., Cherry, C.: SemEval-2016 task 6: Detecting stance in tweets. In: Proceedings of the 10th International Workshop on Semantic Evaluation. pp. 31–41. SemEval '16 (2016)
51. Mukherjee, S., Weikum, G.: Leveraging joint interactions for credibility analysis in news communities. In: Proceedings of the 24th ACM International Conference on Information and Knowledge Management. pp. 353–362. CIKM'15 (2015)
52. Nakov, P., Alam, F., Shaar, S., Da San Martino, G., Zhang, Y.: COVID-19 in Bulgarian social media: Factuality, harmfulness, propaganda, and framing. In: Proceedings of the International Conference on Recent Advances in Natural Language Processing. pp. 997–1009. RANLP '21 (2021)
53. Nakov, P., Alam, F., Shaar, S., Da San Martino, G., Zhang, Y.: A second pandemic? analysis of fake news about COVID-19 vaccines in Qatar. In: Proc. of Conference on Recent Advances in Natural Language Processing. pp. 1010–1021 (2021)
54. Nakov, P., Barrón-Cedeño, A., Elsayed, T., Suwaileh, R., Márquez, L., Zaghouani, W., Gencheva, P., Kyuchukov, S., Da San Martino, G.: Overview of the CLEF-2018 lab on automatic identification and verification of claims in political debates. In: Working Notes of CLEF 2018 – Conference and Labs of the Evaluation Forum. CLEF '18 (2018)
55. Nakov, P., Corney, D., Hasanain, M., Alam, F., Elsayed, T., Barrón-Cedeño, A., Papotti, P., Shaar, S., Da San Martino, G.: Automated fact-checking for assisting human fact-checkers. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence. pp. 4551–4558. IJCAI '21 (2021)
56. Nakov, P., Da San Martino, G., Elsayed, T., Barrón-Cedeño, A., Míguez, R., Shaar, S., Alam, F., Haouari, F., Hasanain, M., Babulkov, N., Nikolov, A., Shahi, G.K., Struß, J.M., Mandl, T.: The CLEF-2021 CheckThat! lab on detecting check-worthy

- claims, previously fact-checked claims, and fake news. In: *Advances in Information Retrieval*. pp. 639–649. Springer International Publishing, Cham (2021)
57. Nakov, P., Giovanni, D.S.M., Elsayed, T., Barrón-Cedeño, A., Míguez, R., Shaar, S., Alam, F., Haouari, F., Hasanain, M., Mansour, W., Hamdan, B., Ali, Z.S., Babulkov, N., Nikolov, A., Shahi, G.K., Struß, J.M., Mandl, T., Kutlu, M., Kartal, Y.S.: Overview of the CLEF-2021 CheckThat! lab on detecting check-worthy claims, previously fact-checked claims, and fake news. LNCS (12880), Springer (2021)
 58. Nakov, P., Márquez, L., Moschitti, A., Magdy, W., Mubarak, H., Freihat, A.A., Glass, J., Randeree, B.: SemEval-2016 Task 3: Community question answering. In: *Proceedings of the 10th International Workshop on Semantic Evaluation*. pp. 525–545. SemEval '15 (2016)
 59. Nguyen, V.H., Sugiyama, K., Nakov, P., Kan, M.Y.: FANG: Leveraging social context for fake news detection using graph representation. In: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. p. 1165–1174. CIKM '20 (2020)
 60. Nikolov, A., Da San Martino, G., Koychev, I., Nakov, P.: Team_Alex at CheckThat! 2020: Identifying check-worthy tweets with transformer models. In: Cappellato et al. [14]
 61. Oshikawa, R., Qian, J., Wang, W.Y.: A survey on natural language processing for fake news detection. In: *Proceedings of the 12th Language Resources and Evaluation Conference*. pp. 6086–6093. LREC '20 (2020)
 62. Pogorelov, K., Schroeder, D.T., Burchard, L., Moe, J., Brenner, S., Filkukova, P., Langguth, J.: FakeNews: Corona virus and 5G conspiracy task at MediaEval 2020. In: *Proceedings of the MediaEval 2020 Workshop*. MediaEval '20 (2020)
 63. Pomerleau, D., Rao, D.: The fake news challenge: Exploring how artificial intelligence technologies could be leveraged to combat fake news <http://www.fakenewschallenge>. (2017)
 64. Popat, K., Mukherjee, S., Strötgen, J., Weikum, G.: Credibility assessment of textual claims on the web. In: *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*. pp. 2173–2178. CIKM '16 (2016)
 65. Pritzkau, A.: NLytics at CheckThat! 2021: Check-worthiness estimation as a regression problem on transformers. In: Faggioli et al. [26]
 66. Sepúlveda-Torres, R., Saquete, E.: GPLSI team at CLEF CheckThat! 2021: Fine-tuning BETO and RoBERTa. In: Faggioli et al. [26]
 67. Shaar, S., Alam, F., Da San Martino, G., Nakov, P.: The role of context in detecting previously fact-checked claims. Arxiv:2104.07423 (2021)
 68. Shaar, S., Alam, F., Da San Martino, G., Nikolov, A., Zaghrouani, W., Nakov, P., Feldman, A.: Findings of the NLP4IF-2021 shared tasks on fighting the COVID-19 infodemic and censorship detection. In: *Proceedings of the Fourth Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda*. pp. 82–92. NLP4IF '21 (2021)
 69. Shaar, S., Alam, F., Martino, G.D.S., Nakov, P.: Assisting the human fact-checkers: Detecting all previously fact-checked claims in a document. arXiv preprint arXiv:2109.07410 (2021)
 70. Shaar, S., Babulkov, N., Da San Martino, G., Nakov, P.: That is a known lie: Detecting previously fact-checked claims. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. pp. 3607–3618. ACL '20 (2020)
 71. Shaar, S., Haouari, F., Mansour, W., Hasanain, M., Babulkov, N., Alam, F., Da San Martino, G., Elsayed, T., Nakov, P.: Overview of the CLEF-2021 Check-

- That! lab task 2 on detecting previously fact-checked claims in tweets and political debates. In: Faggioli et al. [26]
72. Shaar, S., Hasanain, M., Hamdan, B., Ali, Z.S., Haouari, F., Nikolov, A., Kutlu, M., Kartal, Y.S., Alam, F., Da San Martino, G., Barrón-Cedeño, A., Míguez, R., Beltrán, J., Elsayed, T., Nakov, P.: Overview of the CLEF-2021 CheckThat! lab task 1 on check-worthiness estimation in tweets and political debates. In: Faggioli et al. [26]
 73. Shaar, S., Nikolov, A., Babulkov, N., Alam, F., Barrón-Cedeño, A., Elsayed, T., Hasanain, M., Suwaileh, R., Haouari, F., Da San Martino, G., Nakov, P.: Overview of CheckThat! 2020 English: Automatic identification and verification of claims in social media. In: Cappellato et al. [14]
 74. Shahi, G.K.: AMUSED: An annotation framework of multi-modal social media data. arXiv:2010.00502 (2020)
 75. Shahi, G.K., Dirkson, A., Majchrzak, T.A.: An exploratory study of COVID-19 misinformation on Twitter. *Online social networks and media* **22**, 100104 (2021)
 76. Shahi, G.K., Nandini, D.: FakeCovid – a multilingual cross-domain fact check news dataset for COVID-19. In: *Workshop Proceedings of the 14th International AAAI Conference on Web and Social Media* (2020)
 77. Shahi, G.K., Struß, J.M., Mandl, T.: Overview of the CLEF-2021 CheckThat! lab: Task 3 on fake news detection. In: Faggioli et al. [26]
 78. Shu, K., Sliva, A., Wang, S., Tang, J., Liu, H.: Fake news detection on social media: A data mining perspective. *SIGKDD Explor. Newsl.* **19**(1), 22–36 (2017)
 79. Thorne, J., Vlachos, A., Christodoulopoulos, C., Mittal, A.: FEVER: a large-scale dataset for fact extraction and VERification. In: *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. pp. 809–819. NAACL '18 (2018)
 80. Touahri, I., Mazroui, A.: EvolutionTeam at CheckThat! 2020: Integration of linguistic and sentimental features in a fake news detection approach. In: Cappellato et al. [14]
 81. Vasileva, S., Atanasova, P., Màrquez, L., Barrón-Cedeño, A., Nakov, P.: It takes nine to smell a rat: Neural multi-task learning for check-worthiness prediction. In: *Proceedings of the International Conference on Recent Advances in Natural Language Processing*. pp. 1229–1239. RANLP '19 (2019)
 82. Williams, E., Rodrigues, P., Novak, V.: Accenture at CheckThat! 2020: If you say so: Post-hoc fact-checking of claims using transformer-based models. In: Cappellato et al. [14]
 83. Williams, E., Rodrigues, P., Tran, S.: Accenture at CheckThat! 2021: Interesting claim identification and ranking with contextually sensitive lexical training data augmentation. In: Faggioli et al. [26]
 84. Zhao, Z., Resnick, P., Mei, Q.: Enquiring minds: Early detection of rumors in social media from enquiry posts. In: *Proceedings of the 24th International Conference on World Wide Web*. pp. 1395–1405. WWW'15 (2015)
 85. Zhou, X., Wu, B., Fung, P.: Fight for 4230 at CLEF CheckThat! 2021: Domain-specific preprocessing and pretrained model for ranking claims by check-worthiness. In: Faggioli et al. [26]
 86. Zubiaga, A., Liakata, M., Procter, R., Hoi, G.W.S., Tolmie, P.: Analysing how people orient to and spread rumours in social media by looking at conversational threads. *PLoS ONE* **11**(3) (2016)
 87. Zuo, C., Karakas, A., Banerjee, R.: A hybrid recognition system for check-worthy claims using heuristics and supervised learning. In: Cappellato et al. [16]