

Text_Minor at CheckThat! 2022: Fake News Article Detection Using RoBERT.

Sujit Kumar¹, Gaurav Kumar² and Sanasam Ranbir Singh³

Indian Institute of Technology, Guwahati, India

Abstract

Disinformation detection is emerging as an important research challenge due to the rise of disinformation on digital platforms. Several methods have been proposed in the literature to counter the spread of disinformation over digital platforms. However, most of these studies are based on social media, evidence claim verification, and incongruent news article detection. Earlier studies on fake news article detection are based on the stance detection approach over synthetically generated fake news datasets. This paper presents our RoBERT based proposed model submitted to checkThat! task 3 CLEF-2022. We conducted our experiment on the fake news dataset provided by the organizers of task 3 CLEF-2022.

Keywords

Fake news detection, Recurrence over BERT, Misinformation detection,

1. Introduction

The internet and digital platforms have gradually risen as leading sources of news and event information. Studies in literature have revealed various aspects that influence the popularity of social media and digital platforms for news consumption. Compared to conventional newspapers and media, news consumption via social media and online portals is significantly less expensive and early accessible. Although social media and digital platforms provide the easy accessibility of the latest updates to the news consumer, the continuous spread of misinformation such as fake news, clickbait, propaganda, satire or parody and rumors pose a critical threat to the society [1] [2]. Fake news¹ is described as a fabricated storyline on a broad scale to deceive readers. According to media scholars [3], fake news is defined as distorted and deceptive content in circulation as news via communication mediums such as print, electronic, and digital communication.

The first study on fake news detection can be traced back to the year 2014 [4]. Towards the goal to detect fake news in news articles, the first fake news challenge (FNC-1)² was organized by [5] to counter spread of misinformation in form of news article. Several methods have been proposed in the literature for detection of fake news article. This study presents our approach to fake news detection for the shared task at checkThat! for English language dataset. The rest of the paper is organized as follows. Section 2 provides a brief overview on related work,

CLEF 2022: Conference and Labs of the Evaluation Forum, September 5–8, 2022, Bologna, Italy

✉ sujitkumar@iitg.ac.in (S. Kumar); gauravkumar@iitg.ac.in (G. Kumar); ranbir@iitg.ac.in (S. R. Singh)



© 2022 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)

¹https://en.wikipedia.org/wiki/Fake_news

²<http://www.fakenewschallenge.org/>

further section 3 gives the details of shared task and section 4 introduces our proposed model. Section 5 discusses the details about various parameters and hyperparameters used to produce the experimental result. Finally, sections 6 and 7 present the result's analysis of the model and conclusion, respectively.

2. Related Work

In the literature, studies [6], [7], [8], [9], [10], [11], [12], [13], [13], [14] have briefly reviewed and analyzed works related to misinformation and disinformation detection. In this study, we retrospect works related to fake news article detection only. Studies related to fake news article detection can be categorized into three groups: feature-based approach, similarity-based approach and summarization-based approach. Initial studies on fake news article detection utilized bag-of-words-based features for training ensemble models or multi-layer perceptrons. First fake news contest (FNC-1)³ was organized by [5]. The winning system⁴ of Fake News Challenge combined convolutional neural network (CNN) trained over word embeddings of headline and body with Xgboost model trained with bag-of-words based features. Their Xgboost model was trained over count, TF-IDF, SVD, sentiment and word2vec [15] features. The second winner, system Team Athene [16] trained multi-layer perceptron on bag-of-words based and domain-dependent features. Study [17] forms a concatenated feature vector by combining the term frequency-inverse document frequency (TF-IDF) vector of headline and body along with cosine similarity between TF-IDF vector of headline and body. These concatenated features are then used to train a multi-layer perceptron to classify the relationship between the headline and body of a news article. Considering the performance of bag-of-words features-based models in studies⁵, [16] [17] it is evident that bag-of-words based features which include SVD, TF-IDF, count of unigrams, bigrams, trigrams overlap between headline and body features help in fake news article classification. This is not surprising, as bag-of-words feature help capture the similarity between headline and body. However, the feature-based approach [18] fails to consider sequential and contextual information in the headline and body of news articles. The study [18] also suggests that feature-based methods depend upon lexical overlap between headline and body pair. In some cases, though, the headline and body are similar, still the feature-based approach classifies them as unrelated if the body contains a synonym of tokens in the headline rather than tokens. Considering the significance of contextual and sequential information, studies [18] [19] combine bag-of-words-based features with the sequential encoding of headline and body with LSTM [20] and GRU[21]. A news article has a hierarchical structure, where it is defined by a headline and a body. Further, the body is defined by a sequence of paragraphs, and a sequence of sentences define a paragraph. Study [22] explores discourse-level structure between document sentences for fake news detections. Study [23] exploits the hierarchical structure of news article body for incongruent news classifications. The study [23] only considers the hierarchical structure of news articles up to paragraph level. However, the hierarchical structure of news articles can be defined up to the word level. Exploiting the hierarchical structure of

³Fake News Contest-1

⁴First Winner System FNC-1

⁵First Winner System FNC-1

news articles from the body level to word level could help in capturing long-term dependencies between words of a sentence [24] and dependencies between sentences of paragraphs. Here, dependency between words implies that two words may be far away in sentences, but they may be close contextually [24]. Several state-of-the-art document encoders are available in the literature for encoding a sentence by considering long-term dependencies between words, such as tree transformer [25] multiplicative LSTM [26][27]. With the objective to exploit hierarchical structure upto word level for capturing long-term dependencies between words of sentences, we use pretrained BERT⁶[28].

Recent studies [29] [30] applied the summarization technique over a news article body to generate a synthetic headline from it, which represents and summarizes the body. Subsequently, text matching is applied between the generated headline and the actual headline to detect the incongruent headline. Study [31] applied summarization technique which ranks sentences in sentence graph based on the ability of sentences to represent the core concept of the document. The encoding of each sentence in the sentence graph is updated based on its similarity with neighbors. Then, the weighted summation of sentence encoding is passed to a multilayer perceptron for fake news detection. However, synthetically generated headlines from news article body may not be a faithful or good representation of news article body [32], [33]. Suppose the article is partially congruent, with most of the sentences in the body being congruent with the headline except for a few. In that case, a summary of the news article body is dominated by the congruent part of the news article body. Hence, the summarization-based approach fails to detect partially incongruent news articles.

Other sorts of false news, such as *partially false news* items, are still being circulated on social media. To prevent the spread of such false information, study [34] categorized news article in four categories, *fake*, *true*, *partially false* and *other* class. Study [35] created fake news dataset which depicts the genuine characteristics of false news articles that are circulated on social media platforms. Studies [36] [37] present the patterns captured in cross-domain and multilingual fake news detection. The study [36] proposed the first multilingual and cross-domain open-source dataset for misinformation detection during the pandemic time. Study [37] also released large scale multilingual and cross-domain dataset for fake news detection and fact check. Studies [37] proposed a framework to collect and annotate the data. It collects the labelled data from different social media platform in various formats such as image, video or text and annotates the data by a semi-automatic approach.

3. Task Description

CheckThat! contest was organized by Conference and Labs of the Evaluation Forum CLEF2022 to verify the authenticity of news articles. The main objective of shared task 3 [38] was given a pair of text \mathcal{B} and the title \mathcal{H} , classify the title and text in one of the following categories: *true*, *false*, *partially false*, or *other*. If a claim made in a news story is valid, it is said to be *true*. Similarly, a news article is false if the main claim of the news article is *false*. When part of the news article is genuine, and part of the news article is false, then it is classified as *partially false*.

⁶<https://huggingface.co/bert-base-cased>

If a news item does not fit into any category, *true*, *false* or *partially false*, it is placed in *other* class.

4. Proposed system

Inspired by the study [23], we extend the hierarchical structure of news articles from news article body to word level to capture long-term dependencies between words of sentences. Although, several state-of-the-art document encoders are available in literature such as tree transformer [25], multiplicative LSTM [26][27] for encoding a sentence by considering long-term dependencies between words, we considered pretrained BERT⁷[28] for sentence encoding. We did not fine-tune BERT, keeping in mind the limited size of the available training dataset. Ideally, we could have encoded the entire body of the news article using BERT instead of encoding a sentence, but pre-trained BERT does not consider more than 512 tokens [39]. Motivated by such limitations, we proposed Recurrent over BERT (RoBERT) based models. RoBERT captures two significant properties of news article (i) Encoding of a sentence using pre-trained BERT, which captures long-term dependencies between words because of multi-head attention between words in the encoder component of BERT (ii) News article body is a sequence of sentences. We split the news article body into several sentences and applied pre-trained BERT to obtain the encoding of the sentences. Every sentence in the body is related to the previous and next sentences in the news article. Hence, BiLSTM is applied over the encoding of sentences to encode news article body from left to right where every sentence is conditioned over the previous sentence and right to left encoding where every sentence is conditioned over the encoding of the next sentence. Finally, left to right and right to left encoding are concatenated to form the encoding of the news article body.

Figure 1 presents the block diagram of our proposed system. Given news articles \mathcal{N} with text \mathcal{B} and title \mathcal{H} pair. We split text \mathcal{B} into a set of m sentences. We first obtained encoded representation s_i of i^{th} sentence in \mathcal{B} using pretrained Bidirectional Encoder Representations from Transformers (BERT)⁸ [28]. Similarly, we also obtained encoded representation \mathbf{h} of title \mathcal{H} . Then we apply Bidirectional Long short-term memory (BiLSTM) [20] over encoded representation of sentences to obtain the encoded representation \mathbf{b} of text \mathcal{B} . Our system also utilized bag-of-word based features. Utilizing the various features which include overlap features, SVD similarity between text and title, and TF-IDF similarity. As discussed in sections 2 bag-of-words based feature help to capture similarity in terms of lexical overlap. The study [18] suggests that BoW-based features positively impact the fake news detection task. The detail of bag-of-word based features are as follows:

- **Overlap features:** This feature counts several overlaps: unigrams, bigrams and trigrams between text and title. To extract the count overlap feature, first, we extract unigrams, bigrams and trigrams for both text \mathcal{B} and title \mathcal{H} . After that, we counted how many unigrams, bigrams, trigrams of title \mathcal{H} are present in unigrams, bigrams, trigrams of text \mathcal{B} . These features essentially count the common unigrams, bigrams and trigrams between title \mathcal{H} and text \mathcal{B} .

⁷<https://huggingface.co/bert-base-cased>

⁸<https://huggingface.co/bert-base-cased>

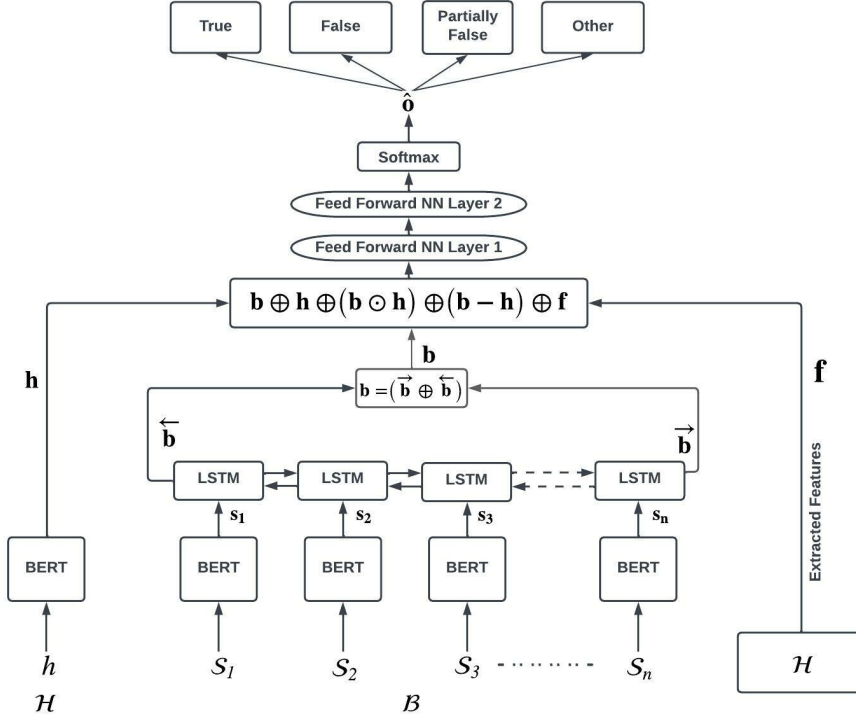


Figure 1: Block diagram of the proposed system. Here \mathcal{S}_i is i^{th} a sentence of text. BERT is applied to obtain encoded representation s_i , \mathbf{h} of text \mathcal{B} , title \mathcal{H} respectively. Then bidirectional LSTM is applied over the encoded representation of sentences to obtain the encoded representation \mathbf{b} of text \mathcal{B} . Finally, feature vectors are estimated to measure the angle and difference between encoded representation \mathbf{b} , \mathbf{h} of text and title, respectively. These estimated features are then passed to a fully connected neural network, followed by Softmax for fake news classification.

- **Singular value decomposition similarity between text and title:** Singular value decomposition (SVD) [40] features help obtain the latent topics involved in the corpus and represent text and title as a mixture of these topics. To obtain SVD of title \mathcal{H} and text \mathcal{B} first, we construct title to words matrix and text to words matrix, where an entry in title to words and text to words matrix are TF-IDF of each word. SVD is then applied over both texts to word and title to word matrix. We retained the top 50 dimensions from both matrices in their decomposition. To obtain similarity between text and title, we apply cosine similarity between SVD of headline and SVD of body.
- **Term Frequency-Inverse Document Frequency (TF-IDF) similarity:** First TF-IDF feature vectors for title and text are obtained by calculating the Term-Frequency of each unigram, normalized by its Inverse-Document Frequency. Then we calculate the cosine similarity between these title and text TF-IDF vectors.

Given encoded representation \mathbf{b} and \mathbf{h} of text \mathcal{B} and title \mathcal{H} respectively. We further obtained

the following feature.

$$\mathbf{r} = \mathbf{b} \odot \mathbf{h} \quad (1)$$

$$\mathbf{d} = \mathbf{b} - \mathbf{h} \quad (2)$$

Now, we define the final feature for the classification as follows.

$$\mathbf{p} = \mathbf{b} \oplus \mathbf{h} \oplus \mathbf{r} \oplus \mathbf{d} \oplus \mathbf{f} \quad (3)$$

where \oplus denotes concatenation and \mathbf{f} is bag-of-words based features. Finally, the estimated feature vector \mathbf{p} is passed through a fully connected neural network followed by Softmax. Our system used cross entropy as loss function to learn the parameters. Our experimental setup was based on 100 LSTM hidden units, two layers fully connected neural network and 500 training epochs.

5. Experimental Setup

This study uses cross-entropy as a loss function with a learning rate of 0.01 to learn the parameters. We consider a maximum of 32 sentences in the text of a news article. If the number of sentences in the text is less than 32, then we pad the random vector, and if there are more than 32 sentences, we consider only the first 32 sentences. Table 1 presents the value of other hyperparameters used for experiment. Our code repository is publicly available at⁹ to reproduce the result presented in this paper: GitHub link https://github.com/SUJIT-KUMAR-ai/Text_Minor-at-CheckThat-2022.

Table 1

Details of hyperparameters used in experimental setup

Hyperparameters	Value
Batch size	4
Learning rate	0.01
Activation function	Softmax
Loss function	Cross entropy
# Epochs	500
LSTM hidden state dimension	100
# Layer in Feedforward NN	2
Max # Sentence in text	32
# BERT dimension	768

6. Result

Table 2 presents the performance of two different setups of the proposed model, with and without features. From Table 2, it can be observed that the performance of the proposed model

⁹Code repository to reproduce results of this paper

is superior by considering bag-of-words-based features. The bag-of-words-based features help the model recognize *other* class samples and boost the performance over *true*, *false* and *partially false* classes. We observed a significant improvement in performance by using bag-of-words-based features. It could be observed from the experiment that the model is performing poorly on the *other* class without using bag-of-words-based features set.

Table 2

Performance Table over Development set and Test set with or without including extracted features

Dataset	Model	Performance					
	RoBERT	Accuracy	F1	true	false	partially false	other
Development Set	with features	0.527	0.502	0.336	0.604	0.515	0.554
	without features	0.406	0.286	0.222	0.504	0.421	0
Test Set	with features	0.442	0.296	0.276	0.619	0.137	0.155
	without features	0.400	0.245	0.227	0.574	0.130	0

7. Conclusion

This paper presents a RoBERT-based model for fake news article detection. We also experimented with other models based on sentence BERT and traditional machine learning. But our RoBERT-based model outperformed other models over the validation set provided by the organizer of checkThat! task 3 CLEF-2022. Accuracy and F1 score of our proposed system submitted to checkThat! task 3 is 0.377 and 0.234, respectively. However, we recreated the experiment with a labeled test dataset released by the organizer checkThat! of task 3 and observed accuracy of 0.442 and an average F1 score of 0.296. For the submission to 'checkThat! Lab Task 3', we have considered the batch parameter as a size of 8 samples. The new results are with a batch size of 4. Hence, there is a slight difference between results. Though, our system performed average compared to other systems submitted to checkThat!, fine-tuning BERT can significantly improve the performance of our proposed method with the large-scale dataset. In future work, we will investigate the performance of our proposed system over the publicly available large-scale dataset.

References

- [1] S. Vosoughi, D. Roy, S. Aral, The spread of true and false news online, *Science* 359 (2018) 1146–1151.
- [2] C. Castillo, M. Mendoza, B. Poblete, Information credibility on twitter, in: *Proceedings of the 20th international conference on World wide web*, 2011, pp. 675–684.
- [3] N. Higdon, *The anatomy of fake news: A critical news literacy education*, University of California Press, 2020.
- [4] K. ARVIND, S. GOVARTHAN, S. K. KUMAR, M. N. KUMAR, R. LAKSHMI, Fake news detection and rumour source identification, *science* 29 (2014) 443–452.

- [5] D. Pomerleau, D. Rao, Fake news challenge, Exploring how artificial intelligence technologies could be leveraged to combat fake news. url: <https://www.fakenewschallenge.org/>(visited on 03/13/2020) (2017).
- [6] K. Shu, A. Sliva, S. Wang, J. Tang, H. Liu, Fake news detection on social media: A data mining perspective, *ACM SIGKDD explorations newsletter* 19 (2017) 22–36.
- [7] S. Kumar, N. Shah, False information on web and social media: A survey, *arXiv preprint arXiv:1804.08559* (2018).
- [8] A. Zubiaga, A. Aker, K. Bontcheva, M. Liakata, R. Procter, Detection and resolution of rumours in social media: A survey, *ACM Computing Surveys (CSUR)* 51 (2018) 1–36.
- [9] K. Sharma, F. Qian, H. Jiang, N. Ruchansky, M. Zhang, Y. Liu, Combating fake news: A survey on identification and mitigation techniques, *ACM Transactions on Intelligent Systems and Technology (TIST)* 10 (2019) 1–42.
- [10] X. Zhou, R. Zafarani, A survey of fake news: Fundamental theories, detection methods, and opportunities, *ACM Computing Surveys (CSUR)* 53 (2020) 1–40.
- [11] S. B. Parikh, P. K. Atrey, Media-rich fake news detection: A survey, in: *2018 IEEE conference on multimedia information processing and retrieval (MIPR)*, 2018, pp. 436–441.
- [12] A. D’Ulizia, M. C. Caschera, F. Ferri, P. Grifoni, Fake news detection: a survey of evaluation datasets, *PeerJ Computer Science* 7 (2021) e518.
- [13] F. Xu, V. S. Sheng, M. Wang, A unified perspective for disinformation detection and truth discovery in social sensing: A survey, *ACM Computing Surveys (CSUR)* 55 (2021) 1–33.
- [14] B. Kim, A. Xiong, D. Lee, K. Han, A systematic review on fake news research through the lens of news creation and consumption: Research efforts, challenges, and future directions, *Plos one* 16 (2021) e0260080.
- [15] T. Mikolov, K. Chen, G. Corrado, J. Dean, Efficient estimation of word representations in vector space, *arXiv preprint arXiv:1301.3781* 1 (2013) 1–12.
- [16] A. Hanselowski, P. Avinesh, B. Schiller, F. Caspelherr, Description of the system developed by team athene in the fnc-1, 2017, Online: https://github.com/hanselowski/athene_system/blob/master/system_description_athene.pdf. Accessed 1 (2018) 03–13.
- [17] B. Riedel, I. Augenstein, G. P. Spithourakis, S. Riedel, A simple but tough-to-beat baseline for the fake news challenge stance detection task, *arXiv preprint arXiv:1707.03264* 1 (2017) 1–6.
- [18] A. Hanselowski, A. PVS, B. Schiller, F. Caspelherr, D. Chaudhuri, C. M. Meyer, I. Gurevych, A retrospective analysis of the fake news challenge stance-detection task, in: *Proceedings of the 27th International Conference on Computational Linguistics*, Association for Computational Linguistics, Santa Fe, New Mexico, USA, 2018, pp. 1859–1874. URL: <https://aclanthology.org/C18-1158>.
- [19] L. Borges, B. Martins, P. Calado, Combining similarity features and deep representation learning for stance detection in the context of checking fake news, *Journal of Data and Information Quality (JDIQ)* 11 (2019) 1–26.
- [20] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural computation* 9 (1997) 1735–1780.
- [21] K. Cho, B. Van Merriënboer, D. Bahdanau, Y. Bengio, On the properties of neural machine translation: Encoder-decoder approaches, *arXiv preprint arXiv:1409.1259* (2014).
- [22] H. Karimi, J. Tang, Learning hierarchical discourse-level structure for fake news detection,

- in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 3432–3442. URL: <https://aclanthology.org/N19-1347>. doi:10.18653/v1/N19-1347.
- [23] S. Yoon, K. Park, J. Shin, H. Lim, S. Won, M. Cha, K. Jung, Detecting incongruity between news headline and body text via a deep hierarchical encoder, Proceedings of the AAAI Conference on Artificial Intelligence 33 (2019) 791–800.
- [24] J. Li, M.-T. Luong, D. Jurafsky, E. Hovy, When are tree structures necessary for deep learning of representations?, arXiv preprint arXiv:1503.00185 1 (2015) 1–11.
- [25] Y. Wang, H.-Y. Lee, Y.-N. Chen, Tree transformer: Integrating tree structures into self-attention, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Association for Computational Linguistics, Hong Kong, China, 2019, pp. 1061–1070.
- [26] N. K. Tran, W. Cheng, Multiplicative tree-structured long short-term memory networks for semantic representations, in: Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics, Association for Computational Linguistics, New Orleans, Louisiana, 2018, pp. 276–286.
- [27] K. S. Tai, R. Socher, C. D. Manning, Improved semantic representations from tree-structured long short-term memory networks, in: Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), Association for Computational Linguistics, Beijing, China, 2015, pp. 1556–1566.
- [28] J. Devlin, M. Chang, K. Lee, K. Toutanova, BERT: pre-training of deep bidirectional transformers for language understanding, CoRR abs/1810.04805 (2018). URL: <http://arxiv.org/abs/1810.04805>. arXiv:1810.04805.
- [29] R. Mishra, P. Yadav, R. Calizzano, M. Leippold, Musem: Detecting incongruent news headlines using mutual attentive semantic matching, in: 2020 19th IEEE International Conference on Machine Learning and Applications (ICMLA), IEEE, 2020, pp. 709–716.
- [30] R. Sepúlveda-Torres, M. Vicente, E. Saquete, E. Lloret, M. Palomar, Headlinestancechecker: Exploiting summarization to detect headline disinformation, Journal of Web Semantics (2021) 100660.
- [31] G. Kim, Y. Ko, Graph-based fake news detection using a summarization technique, in: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, 2021, pp. 3276–3280.
- [32] A. See, P. J. Liu, C. D. Manning, Get to the point: Summarization with pointer-generator networks, in: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Vancouver, Canada, 2017, pp. 1073–1083. URL: <https://aclanthology.org/P17-1099>. doi:10.18653/v1/P17-1099.
- [33] Z. Cao, F. Wei, W. Li, S. Li, Faithful to the original: Fact aware neural abstractive summarization, in: Proceedings of the AAAI Conference on Artificial Intelligence, volume 32, 2018.
- [34] G. K. Shahi, J. M. Struß, T. Mandl, Overview of the clef-2021 checkthat! lab task 3 on fake

- news detection, Working Notes of CLEF (2021).
- [35] G. K. Shahi, A. Dirkson, T. A. Majchrzak, An exploratory study of covid-19 misinformation on twitter, *Online Social Networks and Media* 22 (2021) 100104.
 - [36] G. K. Shahi, D. Nandini, 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. URL: http://workshop-proceedings.icwsm.org/pdf/2020_14.pdf.
 - [37] G. K. Shahi, Amused: An annotation framework of multi-modal social media data, *arXiv preprint arXiv:2010.00502* (2020).
 - [38] J. Köhler, G. K. Shahi, J. M. Struß, M. Wiegand, M. Siegel, T. Mandl, M. Schütz, Overview of the CLEF-2022 CheckThat! lab task 3 on fake news detection, in: *Working Notes of CLEF 2022—Conference and Labs of the Evaluation Forum, CLEF '2022, Bologna, Italy, 2022*.
 - [39] R. Pappagari, P. Zelasko, J. Villalba, Y. Carmiel, N. Dehak, Hierarchical transformers for long document classification, in: *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, IEEE, 2019, pp. 838–844.
 - [40] S. T. Dumais, et al., Latent semantic analysis, *Annu. Rev. Inf. Sci. Technol.* 38 (2004) 188–230.