



Fake News Detection System Using XLNet Model with Topic Distributions: CONSTRAINT@AAAI2021 Shared Task

Akansha Gautam^(✉) , V. Venkatesh , and Sarah Masud 

Indraprastha Institute of Information Technology, Delhi, India
{akansha16221,venkateshv,sarahm}@iiitd.ac.in

Abstract. With the ease of access to information, and its rapid dissemination over the internet (both velocity and volume), it has become challenging to filter out truthful information from fake ones. The research community is now faced with the task of automatic detection of fake news, which carries real-world socio-political impact. One such research contribution came in the form of the Constraint@AAAI2021 Shared Task on COVID19 Fake News Detection in English. In this paper, we shed light on a novel method we proposed as a part of this shared task. Our team introduced an approach to combine topical distributions from Latent Dirichlet Allocation (LDA) with contextualized representations from XLNet. We also compared our method with existing baselines to show that XLNet + Topic Distributions outperforms other approaches by attaining an F1-score of 0.967.

Keywords: Fake news detection · XLNet · LDA · Topic embeddings · Neural network · Natural language processing

1 Introduction

With an increase in the adoption of social media as a source of news, it has become easier for miscreants to share false information with millions of users. Such activities increase during a time of crisis where some groups try to exploit the human vulnerability. One saw during COVID19 the impact of fake news¹ from 5G towers to fad remedies, some even leading to physical harm. Given the volume of fake news generated on a regular basis, there is a need for automated identification of fake news to aid in moderation, as manual identification is cumbersome and time-consuming.

Fake news detection is a challenging problem because of its evolving nature and context-dependent definition of what is fake [1]. For instance, a message shared may have a falsifiable claim but was not shared with the intent to spread misinformation. On the other hand, messages transmitted with the intent to

¹ <https://news.un.org/en/story/2020/04/1061592>.

A. Gautam, V. Venkatesh and S. Masud—Contributed equally.

© Springer Nature Switzerland AG 2021

T. Chakraborty et al. (Eds.): CONSTRAINT 2021, CCIS 1402, pp. 189–200, 2021.

https://doi.org/10.1007/978-3-030-73696-5_18

mislead the masses may contain conspiracy theories. These messages may also include some facts that are not related to the message. While it is relatively easy for a human to identify that the facts mentioned have no relation to the claim made, it is challenging to classify news with such facts as fake automatically. It would require more training samples to induce more discriminatory power in the learned distributed representations. Automatic fake news detection has recently gained interest in the machine learning community. Several methods have been proposed for automatic fake news detection. While initial methods leverage hand-crafted features based on n-grams and psycho-linguistics [15]. Recently, rather than leveraging hand-crafted features, automatic extraction of relevant features in the form of distributed representations has become popular [21]. Various previous studies [6, 7, 20] have shown the effect usage of Language Model Fine-tuning are a better alternative for the classification tasks than other methods.

In this paper, we propose a novel method that combines the contextualized representations from large pre-trained language models like BERT [5] or XLNet [23] with Topic distributions from Latent Dirichlet Allocation (LDA) [3] for the COVID19 Fake News Detection in English competition [13]. We observed that the topic distribution provides more discriminatory power to the model. The joint representations are used for classifying the inputs as ‘fake’ or ‘real’. Since the given shared task contains domain-specific language, we posit that topic distributions help provide additional signals that improve overall performance. The topic models have been previously exploited for domain adaptation [8].

Our core technical contributions are in four areas:

- We propose a novel system for fake news detection that combined topic information and contextualized representations (Sect. 3).
- We provide an extensive comparison with other states of art the neural methods and rudimentary machine learning models (Sect. 5.2).
- We attempt to perform error analysis both in terms of term-token counts and attention heads (Sect. 6).
- We provide the source code² use for modeling and error analysis along with values of hyper-parameters (Sect. 5.1).

2 Related Work

Several researchers have already contributed by designing a novel approach to solving the problem of automatic fake news detection. A group of researchers [15] developed two datasets named *Celebrity* and *FakeNewsAMT* that contains equal proportions of real and fake news articles. They use linguistic properties such as n-grams, punctuation, psycho-linguistic features, readability, and Syntax to identify fake articles. They use linear SVM classifier as a baseline model to conduct several experiments such as learning curve and cross-domain analyses with a different combination of features set.

² Source code available at: <https://github.com/VenkteshV/Constraint2021>.

Another group of researchers [17] identified the characteristics of fake news articles into three parts: (1) textual data of article (2) response of user (3) source users promoting articles. They proposed a model called CSI composed of Capture, Score, and Integrate modules. The first module uses the Recurrent Neural Network (RNN) to capture the temporal representations of articles. The second module is based on the behavior of users. The third module uses the output produced by the first two models to identify fake news articles. Some prior studies [19] have also used news content with additional information (social context information) to build a model to detect fake news. In a parallel study, [11] of fake news in China, the hybrid model assimilates the speaker profile into the LSTM. The research shows that speaker profiles help in improving the Fake News Detection model's accuracy.

A study [21] used LIAR dataset. They proposed a model based on surface-level linguistic patterns. The baseline includes logistic regression, support vector machines, long short-term memory networks, and a convolutional neural networks model. They designed a novel, hybrid convolutional neural network to integrate metadata with text, which achieved significant fine-grained fake news detection.

A group of researchers [2] presented a robust and straightforward model for the Shared Task on profiling fake news spreaders. Their method relies on semantics, word classes, and some other simple features and then fed these features to the Random Forest model to classify fake news. The study [10] focuses on introducing a novel method for detecting fake news on social media platforms. They used news propagation paths with both recurrent and convolutional networks to capture global and local user characteristics.

A recent study [16] presented a new set of features extracted from news content, news source, environment. It measured the prediction performance of the current approaches and features for the automatic detection of fake news. They have used several classic and state-of-the-art classifiers, including k-Nearest Neighbors, Naive Bayes, Random Forest, Support Vector Machine with RBF kernel, and XGBoost to evaluate the discriminative power of the newly created features. They measure each classifier's effectiveness with respect to the area under the ROC curve (AUC) and the Macro F1-score. Another recent study [18] focuses on two variations of end to end deep neural architectures to identify fake news in the multi-domain platform. The first model is based on Bidirectional Gated Recurrent Unit (BiGRU) comprised of (1) Embedding Layer (2) Encoding Layer (Bi-GRU) (3) Word-level Attention (4) Multi-layer Perceptron (MLP). However, another model is based on Embedding from Language Model (ELMo) and the MLP Network.

3 Proposed Method

This section describes in detail the proposed approach. The proposed neural network architecture for the fake news detection task is shown in Fig. 1. We leverage contextualized representations from XLNet and representations obtained from




Fig. 1. Proposed Model Architecture using XLNet with Topic Distributions, where contextualized representations and topic embeddings are obtained from the XLNet and LDA model, respectively. These representations are then concatenated and fed to the 2-fully connected layer followed by a Softmax Layer for the task of fake news detection.

Latent Dirichlet Allocation (LDA) to obtain useful representations for fake news classification. The LDA is a generative probabilistic model. Each word in the document d is assumed to be generated by sampling a topic from d 's topic distribution θ^d and then sampling a word from the distribution over words denoted by ϕ^t of a topic.

We leverage contextualized representations to handle the problem of polysemy. The problem of polysemy occurs when the same word has different meanings in different contexts. The vector representations obtained through methods like word2vec are unable to disambiguate such terms and hence output the exact representations for the word irrespective of the context of their occurrence. The recent wave of pre-trained language models is based on the transformer architecture, which uses a mechanism called self-attention. The self-attention mechanism computes better representations for a word in a sentence by scoring other words in the sentence against the current word. This helps determine the amount of focus placed on different input sequence words when computing the present word's vector representation. The pre-trained language model BERT [5] was built using the transformer and provided useful contextualized representations for many downstream tasks. However, there were several drawbacks to BERT. During training, the BERT model predicts the masked tokens in parallel. This may result in wrong predictions as the value of one of the masked tokens may depend on another masked token. For instance, for the sentence “I went to the [MASK] to get [MASK]”. Here, the words “hospital” and “vaccinated” for the first and second masks are more probable than the words “hospital” and “coffee”. However, there are many possible combinations when the BERT model predicts the tokens in parallel, resulting in an incorrect sentence. The XLNet model [23] helps overcome certain drawbacks of BERT by introducing the *permutation language modeling* and by using transformer-XL architecture [4] as the backbone. The transformer-XL architecture introduces the recurrence mechanism at the segment level to the transformer architecture. It accomplished this by caching the hidden states generated from the previous segment and uses them as keys and values when processing the next segment. The permutation language mod-

eling method predicts one token at a time given the preceding context like a traditional language model. However, it predicts the tokens at random order rather than the sequential one. Hence, the permutation language modeling does not need the [MASK] tokens and does not have independent parallel predictions observed in BERT.

In the proposed method, the news article (denoted as a_i) is passed through XLNet model to obtain contextualized representations (denoted as $CE(\cdot)$). The LDA model is trained on the provided training set and is then leveraged to compute the document-topic embeddings (denoted as $TE(\cdot)$) for a given input. The training of LDA is done only once and hence does not add to the inference time. The document-topic distributions can be pre-computed for the entire training set. This can be accomplished easily in our architecture as the computation of the document-topic distributions is decoupled from the XLNet forward pass. The final input representation can be obtained by combining the input's topic distribution with the contextualized embeddings of the sentence. We denote the final input representation as IE , as shown below:

$$IE(a_i) = [[CE(t), TE(t)] | t \in a_i] \quad (1)$$

The concatenated feature representation is passed through 2-fully connected layers followed by a Softmax Layer to output the prediction y_i for classification of news articles.

$$y_i = \text{Softmax}(IE(a_i)) \quad (2)$$

We perform extensive experiments by varying the model architecture. The dataset, experiments conducted, and the baselines are discussed in the following section.

4 Dataset Description

We use the COVID-19 Fake News Dataset given by [14]. It is a manually annotated dataset of 10,700 social media posts and articles of real and fake news based on topics related to COVID-19. Fake news articles are collected from several fact-checking websites and tools, whereas real news articles are collected from Twitter using verified Twitter handles. Table 1 depicts examples of fake and real articles from the COVID-19 Fake News Dataset.

The dataset is split into 6420 samples in the train set and test and validation sets with 2140 samples. Table 2 shows the distribution of data across 2 different classes. It suggests that data is class-wise balanced and class distribution is similar across Train, Validation, and Test Split.

Table 1. Examples of fake and real news articles from the dataset

Label	Text
Fake	No Nobel Prize laureate Tasuku Honjo didn't say the coronavirus is "not natural" as a post on Facebook claims. In fact Professor Honjo said he's "greatly saddened" his name was used to spread misinformation. This and more in the latest #CoronaCheck: https://t.co/rLcTuIcIHO https://t.co/WdoocCiXFu
Real	We launched the #COVID19 Solidarity Response Fund which has so far mobilized \$225+M from more than 563000 individuals companies & philanthropies. In addition we mobilized \$1+ billion from Member States & other generous to support countries-@DrTedros https://t.co/xgPkPdv0r

Table 2. Distribution of dataset across 2 different classes, Real and Fake

Split	Real	Fake	Total
Train	3360	3060	6420
Validation	1120	1020	2140
Test	1120	1020	2140
Total	5600	5100	10700

5 Experiments

The proposed approach was implemented using PyTorch and with an NVIDIA Tesla K80 GPU. We use *Transformers* library³ maintained by the researchers and engineers at Hugging Face [22] which provides PyTorch interface for XLNet. *Transformers* library supports Transformer-based architectures such as BERT, RoBERTa, DistilBERT, XLNet [23] and facilitates the distribution of pre-trained models.

5.1 Implementation

The pre-processing of data involves in our approach is inspired from various sources [9, 14]. We pre-processed the data by removing emojis, stopwords, special characters, hashtag symbols, usernames, links, and lowercasing the texts. We use `xlnet-base-cased` model to conduct our experiment. To provide input to the XLNet model, we first split our text into tokens and mapped these tokens to their index in the tokenizer vocabulary. For each tokenized input text, we construct the following:

- **input ids:** a sequence of integers identifying each input token to its index number in the XLNet tokenizer vocabulary

³ <https://github.com/huggingface/transformers>.

- **attention mask**: a sequence of 1s and 0s, with 1s for all input tokens and 0s for all padding tokens
- **topic embeddings**: a sequence of probabilities signifies the likelihood of a word in conjunction with a given topic using LDA model
- **labels**: a single value of 1 or 0. In our task, 1 means “Real News,” and 0 means “Fake News.”

The model is fine-tuned for 15 epochs with a learning rate of $2e^{-5}$ and an epsilon value of $1e^{-8}$.

5.2 Comparison with Other Methods

We compare our results with the baseline [14] method and our other experimented methods. The explanation about our other methods are mentioned as follows:

- **USE + SVM**: We first adopt a ML-based approach. Instead of TF-IDF features, we obtain contextualized representations of the input using Universal sentence encoder (USE)⁴. We then fed the input representations to an SVM model.
- **BERT with Topic Distributions**: In this approach, we combine the document-topic distributions from LDA with contextualized representations from BERT. The model was fine-tuned for 10 epochs (with early stopping) with the ADAM optimizer, with a learning rate of $2e^{-5}$ and epsilon is set to $1e^{-8}$.
- **XLNet**: Here, we fine-tune the pre-trained XLNet model on the given input. This model was fine-tuned for 15 epochs with ADAM optimizer using the learning rate of $2e^{-5}$, and epsilon is set to $1e^{-8}$.
- **Ensemble Approach: BERT and BERT + topic**: Here, we combine the predictions of the BERT and BERT + topic models. This provides an increase in performance on the validation set. However, this variant does not outperform the proposed XLNet with the Topic Distributions model on the test set.

Table 3 shows the performance of baseline, experimented, and final proposed method using several evaluation metrics such as Precision, Recall, and weighted F1-score on the Test set. It suggests that our proposed method outperforms the baseline and other models by achieving an F1-score of 0.967.

5.3 Results and Discussion

The results of the comparison of the proposed method with baselines and several variants of the proposed method are shown in Table 3. From the table, it is evident that including topical information enhances the performance as BERT + topic outperforms the baseline methods and is similar to the performance


⁴ <https://tfhub.dev/google/universal-sentence-encoder-large/3>.

Table 3. Performance comparison of proposed method with baseline and other variants on Test set

Method	Precision	Recall	F1-score
Baseline method [14]	0.935	0.935	0.935
USE + SVM	0.92	0.92	0.92
BERT with topic distributions	0.949	0.948	0.948
XLNet	0.949	0.948	0.948
Ensemble approach: BERT and BERT + topic	0.966	0.966	0.966
XLNet with topic distributions (Proposed method)	0.968	0.967	0.967

of XLNet. Also, XLNet with Topic Distributions outperforms all methods. We also observe that the difference in F1 scores between the ensemble approach and XLNet with Topic Distributions is not statistically significant. The above results support the hypothesis that topic distributions help in domain adaptation enhancing performance. The topic distributions are pre-computed and hence can be indexed, making our method efficient for inference.

6 Error Analysis

**Fig. 2.** Confusion Matrix of proposed method on Test Set

Based on the Fig. 2, we see that there are a total of 69 misclassified samples. Let us look at a few of these misclassified test samples based on how the sample keywords are distributed across the fake and real classes in the training + validation set combined.

- **EXAMPLE 1 (Test ID 351, Real Classified as Fake):** *today there are 10 people in hospital who have covid-19 three people are in auckland city hospital four people in middlemore two people in north shore hospital and one person in waikato hospital he new person in auckland city hospital is linked to the community cluster.* As we observe from Table 4a that the combined negative




Fig. 3. Attention weights for the terms “chronic” and “lasting” for attention head 7 at layer 7 of fine-tuned XLNet. Input is Example 6.

impact of terms “covid-19” and “hospital” is much greater than the positive impact of the term “people”, which could explain why the prediction skews towards the “Fake” class instead of its actual “Real” class.

- **EXAMPLE 2 (Test ID 186, Real Classified as Fake):** *the claim stated that india’s top business conglomerate tata group chairman ratan tata said it’s not time to think of profits but to think of survival.* Similar to previous example we observe (Table 4b) that the negative impact of term “claim” is much greater than the positive impact of the word “india’s”, which again causes the prediction to skew towards the “Fake” class instead.
- **EXAMPLE 3 (Test ID 1145, Fake Classified as Real):** *there are 59 positive coronavirus cases in nagpur along with three doctors, one of whom is on ventilator.* As we see from Table 4c, the positive impact of the terms “positive”, “cases” and “ventilator”, outweigh the negative impact of the term “coronavirus”. Now, had XLNet attention given more weightage to the negative term “coronavirus”, the predictions would have been on point, but that does not seem to be happening for this example.
- **EXAMPLE 4 (Test ID 468, Fake Classified as Real):** *millions of app users’ send in 3900 photo’s of 8 different type of rashes, so now they’re a covid19 symptom www.* As we observe from Table 4d, the minor negative impact of the term “million” is matched by the minor positive impact of the terms “different” and “symptoms”. Meanwhile the seemingly important keyword “rashes” is not observed at all in any of the training samples. It is however, the highly positive impact of the term “covid19” that skews the prediction in favour of class “Real” instead of “Fake”.
- **EXAMPLE 5 (Test ID 1147, Fake Classified as Real):** *these people have been assessed as presenting a very low risk due to the nature of their exemption adherence to their required protocols and the negative test results of people associated with their bubble.* We see that unfortunately all keywords are contributing positively, giving way to the prediction being “Real” rather than “Fake”.

Table 4. Word token occurrence of some keyword in the above examples. The count is based on the combined training and validation set for the two classes.

Keyword	Count Fake Class	Count Real Class
People	358	581
Hospital	212	141
Covid-19	1194	880

(a) Keyword occurrence of most contributing words in Example 1.

Keyword	Count Fake Class	Count Real Class
Positive	128	212
Coronavirus	1590	371
Cases	194	2003
Ventilator	8	10
Doctors	–	–

(c) Keyword occurrence of most contributing words in Example 3.

Keyword	Count Fake Class	Count Real Class
Claim	139	1
India’s	11	59
Tata	5	1
Survival	6	1

(b) Keyword occurrence of most contributing words in Example 2.

Keyword	Count Fake Class	Count Real Class
Millions	22	6
Different	22	39
Rashes	–	–
Covid19	255	1545
Symptom	3	11

(d) Keyword occurrence of most contributing words in Example 4.

Keyword	Count Fake Class	Count Real Class
People	358	581
Low	15	83
Risk	25	183
Negative	16	80
Test	97	222
Results	17	67

(e) Keyword occurrence of most contributing words in Example 5.

Keyword	Count Fake Class	Count Real Class
Chronic	3	11
Covid-19	1194	880
Health	153	370
Effects	12	12

(f) Keyword occurrence of most contributing words in Example 6.

- **EXAMPLE 6 (Test ID 663, Real Classified as Fake):** *chronic covid-19 has long-lasting health effects.* We see here that the while the while combine impact of the terms “covid-19” and “health” is tilted towards positive, the predicted output is “Fake”. Since, this result cannot be directly explained in terms of term count, we dig deeper and found that the overall attention given to the term “covid” is higher than that of the term “health”. For 7th attention head, of the 7th layer (3), un-normalised attention weight for term “covid” is around ≈ 1.0 , while that of “health” and “effects” combined lags at ≈ 0.3 . This difference in attention weights and the skewed class-wise count have the combined affect of shifting the predicted distribution towards “Fake”.

Some of the techniques that can help reduce this bias towards count could be inclusion of theme specific stop words (common terms like doctors, tests which are related to Covid-19), weighing token weights in XLNet by tf-idf based techniques (give importance of rare words in each class), manual mapping of abbreviations and similar terms to consolidate their impact (“covid19, covid-19, coronavirus” all point to same entity).

7 Conclusion

This paper has proposed a fake news detection system that exploits transfer learning with the LDA model. We used the XLNet model with Topic Distributions derived from the LDA model. Our proposed approach provides a gain in performance when compared to other neural models. We attribute the gain to the topic distributions, which provide more discriminatory power. In the future, we aim to leverage BERTweet [12], which has been trained on a corpus of 850M English tweets. The tweets are different from traditional text from Wikipedia in terms of sequence length and frequent use of informal language. Hence, a language model pre-trained on a large corpus of tweets would help increase the performance when fine-tuned for domain specific tasks.

References

1. Adriani, R.: The evolution of fake news and the abuse of emerging technologies. *Eur. J. Soc. Sci.* **2**(1), 32–38 (2019). <https://doi.org/10.26417/ejss-2019.v2i1-53>, <http://journals.euser.org/index.php/ejss/article/view/4241>
2. Agirrezabal, M.: KU-CST at the profiling fake news spreaders shared task. In: CLEF (2020)
3. Blei, D.M., Ng, A.Y., Jordan, M.I.: Latent dirichlet allocation. *J. Mach. Learn. Res.* **3**(null), 993–1022 (2003)
4. Dai, Z., Yang, Z., Yang, Y., Carbonell, J., Le, Q., Salakhutdinov, R.: Transformer-XL: attentive language models beyond a fixed-length context. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics, Florence, Italy, July 2019. <https://www.aclweb.org/anthology/P19-1285>
5. Devlin, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. *CoRR abs/1810.04805* (2018). <http://arxiv.org/abs/1810.04805>
6. Eisenschlos, J., Ruder, S., Czapla, P., Kardas, M., Gugger, S., Howard, J.: Multifit: efficient multi-lingual language model fine-tuning. *arXiv preprint arXiv:1909.04761* (2019)
7. Howard, J., Ruder, S.: Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146* (2018)
8. Hu, Y., Zhai, K., Eidelman, V., Boyd-Graber, J.: Polylingual tree-based topic models for translation domain adaptation. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1166–1176. Association for Computational Linguistics, Baltimore, Maryland, June 2014. <https://doi.org/10.3115/v1/P14-1110>, <https://www.aclweb.org/anthology/P14-1110>
9. Kumar, P., Singh, A.: NutCracker at WNUT-2020 task 2: robustly identifying informative covid-19 tweets using ensembling and adversarial training. *arXiv preprint arXiv:2010.04335* (2020)
10. Liu, Y., Wu, Y.F.B.: Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In: Thirty-second AAAI conference on artificial intelligence (2018)

11. Long, Y., Lu, Q., Xiang, R., Li, M., Huang, C.R.: Fake news detection through multi-perspective speaker profiles. In: Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers), pp. 252–256. Asian Federation of Natural Language Processing, Taipei, Taiwan, November 2017. <https://www.aclweb.org/anthology/I17-2043>
12. Nguyen, D.Q., Vu, T., Nguyen, A.T.: BERTweet: a pre-trained language model for English tweets. arXiv preprint [arXiv:2005.10200](https://arxiv.org/abs/2005.10200) (2020)
13. Patwa, P., et al.: Overview of constraint 2021 shared tasks: detecting English covid-19 fake news and Hindi hostile posts. In: Chakraborty, T., Shu, K., Bernard, R., Liu, H., Akhtar, M.S. (eds.) CONSTRAINT 2021, CCIS, vol. 1402, pp. 42–53. Springer, Cham (2021)
14. Patwa, P., et al.: Fighting an infodemic: Covid-19 fake news dataset (2020)
15. Pérez-Rosas, V., Kleinberg, B., Lefevre, A., Mihalcea, R.: Automatic detection of fake news. In: Proceedings of the 27th International Conference on Computational Linguistics, pp. 3391–3401. Association for Computational Linguistics, Santa Fe, New Mexico, USA, August 2018. <https://www.aclweb.org/anthology/C18-1287>
16. Reis, J.C.S., Correia, A., Murai, F., Veloso, A., Benevenuto, F.: Supervised learning for fake news detection. *IEEE Intell. Syst.* **34**(2), 76–81 (2019). <https://doi.org/10.1109/MIS.2019.2899143>
17. Ruchansky, N., Seo, S., Liu, Y.: CSI: a hybrid deep model for fake news detection. In: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, pp. 797–806 (2017)
18. Saikh, T., De, A., Ekbal, A., Bhattacharyya, P.: A deep learning approach for automatic detection of fake news. arXiv preprint [arXiv:2005.04938](https://arxiv.org/abs/2005.04938) (2020)
19. Shu, K., Sliva, A., Wang, S., Tang, J., Liu, H.: Fake news detection on social media: a data mining perspective. *ACM SIGKDD Explor. Newslett.* **19**(1), 22–36 (2017)
20. Sun, C., Qiu, X., Xu, Y., Huang, X.: How to fine-tune BERT for text classification? In: Sun, M., Huang, X., Ji, H., Liu, Z., Liu, Y. (eds.) CCL 2019. LNCS (LNAI), vol. 11856, pp. 194–206. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-32381-3_16
21. Wang, W.Y.: “liar, liar pants on fire”: a new benchmark dataset for fake news detection. arXiv preprint [arXiv:1705.00648](https://arxiv.org/abs/1705.00648) (2017)
22. Wolf, T., et al.: Huggingface’s transformers: state-of-the-art natural language processing. arXiv preprint [arXiv:1910.03771](https://arxiv.org/abs/1910.03771) (2019)
23. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R.R., Le, Q.V.: XLNet: generalized autoregressive pretraining for language understanding. In: Advances in Neural Information Processing Systems, pp. 5753–5763 (2019)