Fake News Stance Detection Using DeepLearning Architecture (CNN)

[1]SWAMINATHAN B, [2]VINISHA K, [3]JEYA BRUNDHA K

BANNARI AMMAN INSTITUTE OF TECHNOLOGY

ABSTRACT

In the contemporary landscape of information dissemination, the proliferation of fake news poses a significant challenge, impacting societal discourse and decision-making processes. Leveraging advancements in deep learning techniques, particularly Convolutional Neural Networks (CNNs), has emerged as a promising approach for discerning the authenticity of news content. This paper introduces a novel two-phase benchmark model, termed WELFake, designed for fake news detection. The first phase involves preprocessing the dataset and validating news content veracity through linguistic features, while the second phase integrates linguistic feature sets with word embedding (WE) and employs a voting classification scheme. To evaluate the efficacy of our approach, we meticulously curate the WELFake dataset comprising approximately 72,000 articles, amalgamating various datasets to ensure unbiased classification outcomes. Experimental results demonstrate that the WELFake model achieves a remarkable classification accuracy of 96.73%, representing a significant enhancement over existing methodologies. Specifically, our model surpasses the accuracy of bidirectional encoder representations from transformer (BERT) by 1.31% and outperforms CNN models by 4.25%. Moreover, comparative analysis with predictive-based approaches utilizing the Word2vec word embedding method showcases an improvement of up to 1.73%. The findings underscore the effectiveness of our frequency-based and focused analysis of writing patterns, affirming the utility of the WELFake model in combating the dissemination of fake news in real-time social media environments. Additionally, the model achieves an overall accuracy of 93%.

KEYWORDS : *Fake news detection, text mining, deep learning, Classification, Benchmark model, WELFake dataset, CNN, wordembedding, Bidirectional encoder representations from transformer (BERT), convolutional neural network (CNN), Word2vec and Social media*

I. INTRODUCTION

In recent years, the unrestricted ability of users to disseminate information on online news platforms, particularly social media, has led to the widespread propagation of misleading and false information [1]. Platforms such as Twitter, Facebook, Instagram, and YouTube have become primary sources of news consumption globally, especially in developing nations, enabling individuals to publish statements and distribute fake news with ease [2].

The ramifications of this phenomenon extend across various sectors of society, business, and culture, with both detrimental and beneficial impacts [2]. Fake news, in particular, has emerged as a significant threat to global commerce, journalism, and democracy, causing substantial collateral damage. For instance, false reports, such as the one claiming that US President Barack Obama had been injured in an explosion, led to a staggering $130 billion loss in the stock market [3]. Studies reveal that a significant portion of fake news originates from official news outlets and online social media platforms, amplifying its societal impact [4]. Given the severity of its consequences, combating fake news through effective detection systems is imperative.

Advancements in Artificial Intelligence (AI) have spurred research into addressing novel challenges, including fake news detection [5]-[8]. Machine Learning (ML) approaches, in particular, have been extensively studied to combat the emergence and dissemination of false news. These systems assist consumers in filtering content and discerning the authenticity of news pieces [5], [9]. The recent successes of Deep Learning (DL) techniques, particularly in natural language processing tasks, have rendered them viable for detecting fake news efficiently and accurately [6]-[8].

In response to the escalating issue of fake news, considerable efforts have been directed toward developing automatic detection systems, leveraging advancements in Artificial Intelligence (AI) and Machine Learning (ML) [5]-[8]. These systems aim to assist consumers in discerning the authenticity of news content, particularly on social media platforms, where misinformation is rampant. Deep Learning (DL) techniques, with their recent successes in natural language processing tasks, offer promising avenues for effectively and efficiently detecting fake news [6]-[8].

Among ML approaches, Convolutional Neural Networks (CNNs) have garnered attention for their effectiveness in various fields, including natural language processing. CNNs utilize convolution layers, pooling layers, and fully connected layers to extract features from datasets, making them well-suited for fake news detection tasks. In this context, we propose an Optimal CNN model for Fake news detection (OPCNN-Fake), designed to extract both high-level and low-level features from news datasets, demonstrating superior performance compared to alternative models.

The urgent need for solutions to verify the authenticity of online content underscores the importance of ongoing research in fake news detection. While ML models with different feature sets have been developed for this purpose, several key questions remain unanswered:

Which linguistic features are most effective in classifying news data as real or fake?
How do different word embedding (WE) techniques, coupled with linguistic features, compare to ML methods like CNNs and Bidirectional Encoder Representations from Transformers (BERTs)?
Which classification method is most suitable for fake news detection across available datasets?
Can ensemble voting classifiers enhance the accuracy of fake news detection systems?
To address these questions, we introduce a novel method named WELFake, which focuses exclusively on text data and operates in three stages: (16) prediction of fake news using linguistic feature sets, (17) enhancement of fake news detection through WE over linguistic feature sets using the WELFake dataset, and (18) comparative analysis of results with state-of-the-art CNN and BERT methods. The WELFake model does not rely on additional metadata information for classification, underscoring its simplicity and effectiveness in detecting fake news on social media platforms.

This paper presents the WELFake model's design, implementation, and evaluation on a comprehensive dataset. Our experimental results demonstrate the model's superior performance, achieving a fake news classification accuracy of up to 96.73%, thereby outperforming existing CNN and BERT methods. The subsequent sections of this paper delve into the methodology, results, discussion, and conclusions, providing insights into the efficacy of the WELFake model and avenues for future research and development.

ll. Related works:

Ahmed et al. [21]: Ahmed et al. conducted experiments on fake news detection using the Kaggle-EXT dataset, comprising 25,200 articles. Their approach did not incorporate linguistic features; instead, they employed a linear Support Vector Machine (SVM) model on TF-IDF representations, achieving an accuracy of 92%.

Shu et al. [22]: Shu et al. focused on fake news detection using linguistic features on the BuzzFeed and Politifact datasets, consisting of 240 and 182 articles, respectively. They utilized a linguistic feature method with two features and implemented SVM classifiers separately on each dataset. Their model achieved accuracies of 87.8% for Politifact and 86.4% for BuzzFeed.

Gravanis et al. [23]: Gravanis et al. employed the UNBiased dataset, which comprised 3,404 articles, including 2,004 real news articles and 1,400 fake ones. They achieved an accuracy of up to 95% using an SVM classifier and utilized 57 linguistic features with the Word2vec word embedding (WE) method. Additionally, they compared their approach on other datasets, such as Kaggle-EXT, BuzzFeed, Politifact, and McIntire, achieving accuracies of up to 99.0%, 72.70%, 84.7%, and 81%, respectively. Notably, they highlighted the use of a biased Kaggle-EXT dataset, containing real news from a single source in 2018, which consistently produced high performance regardless of the number of features.

WELFake: Our proposed WELFake model is applied on a larger dataset comprising over 72,000 news articles, achieving a higher accuracy of 96.73% compared to the aforementioned models. To ensure a fair comparison, we separately evaluated the WELFake model on the four smaller datasets, including Kaggle, McIntire, Reuters, and Buzzfeed. Our results demonstrate significant improvements over Gravanis et al.'s method, with accuracy enhancements from 72.7% to 82.7% on the Buzzfeed dataset and from 81% to 91.78% on the McIntire dataset[22][23].

In [24], a deep learning method is used for addressing
the stance detection problem from the FNC-1 task. It incorporates bi-directional RNNs together with max-pooling and neural attention mechanisms to build representations from headlines and from the body of news articles and combine these representations with external similarity features.

Ahmed et al. [25]: Ahmed et al. utilized N-gram and TF-IDF methods for feature extraction and employed classifiers such as Stochastic Gradient Descent (SGD), Support Vector Machine (SVM), Linear SVM, K-Nearest Neighbor (KNN), and Decision Tree (DT) on the ISOT Fake News dataset. Their approach achieved an accuracy of 92% using the Linear SVM classifier.

Ozbay and Alatas [26]: Ozbay and Alatas employed TF-IDF as the sole feature extraction method and experimented with 23 classifiers, including ZeroR, DT, and Weighted Instances Handler Wrapper (WIHW). They reported superior results compared to Ahmed et al. [13], achieving an accuracy of 96.8% along with high precision, recall, and F1-scores.

Ahmad et al. [27]: Ahmad et al. compared individual learning algorithms such as Logistic Regression (LR), Linear SVM, Multilayer Perceptron (MLP), and KNN with ensemble learning techniques like Random Forest (RF), Voting Classifier, Bagging Classifier, and Boosting Classifier. Their experiments across multiple datasets demonstrated the efficacy of RF algorithm, achieving impressive accuracy, precision, recall, and F1-scores, surpassing previous studies.

Kaliyar et al. [28]: Kaliyar et al. proposed a novel approach by combining pre-trained word embeddings (GloVe) with Convolutional Neural Network (CNN) architecture for fake news detection using the Fake News Dataset. Their method outperformed previous studies, achieving high accuracy, precision, recall, and F1-scores.

Bahad et al. [29]: Bahad et al. also utilized GloVe pre-trained word embeddings and explored various deep learning architectures such as CNN, RNN, Unidirectional LSTM, and Bidirectional LSTM. Their experiments yielded varied results, with some models achieving superior performance compared to previous studies on different datasets.

Deepak & Chitturi [30]: Deepak & Chitturi investigated the impact of secondary features such as news domains, writers, and headlines on fake news detection. They combined word embeddings (BoW, Word2Vec, GloVe) with Feed-forward Neural Network (FNN) and LSTM architectures. Despite observing a significant increase in performance, particularly with LSTM, their results did not surpass the performance reported in some previous studies.

Stance detection, a fundamental task in natural language processing (NLP), involves determining the attitude of a text towards a specific target, such as favor, against, or neutral [31]. This section presents a comprehensive review of existing literature on stance detection, particularly in the context of fake news detection.

Previous studies have established stance detection as a cornerstone for various tasks, including fake news detection, claim validation, and argument search [25], [26], [32]. Target-specific stance prediction has been a focus in many research efforts, especially in analyzing tweets and online debates [33], [34], [28]. These approaches often leverage structural, linguistic, and lexical features to predict stances towards specific claims or topics.

In the realm of fake news detection, stance detection plays a crucial role in determining the alignment of news articles with their headlines [45]. The Fake News Challenge (FNC-1) sparked significant interest in this area, leading to the development of innovative models for stance detection. The top-performing systems in

FNC-1 utilized ensemble methods, deep convolutional neural networks (CNN), and gradient-boosted decision trees, achieving accuracies of up to 82.02% [46].

Subsequent research efforts have explored various approaches to stance detection, including the integration of hand-crafted features, neural network architectures, and attention mechanisms [35]. Deep learning methods, such as bi-directional recurrent neural networks (RNNs) and attention mechanisms, have shown promising results in capturing semantic representations from headlines and news articles, [30]. Additionally, the incorporation of global features and local word embeddings has led to improvements in stance prediction accuracy.

Recent advancements have introduced novel architectures, such as stacked Bi-LSTM and CNN layers, to enhance the modeling of textual sequences. These approaches leverage the strengths of both CNN and LSTM layers for more robust stance detection. Furthermore, large-scale language models, such as Roberta, have demonstrated impressive performance through transfer learning on benchmark datasets [36].

However, existing machine learning models often rely on hand-crafted features, limiting their ability to capture contextual information effectively. Moreover, many models struggle to achieve satisfactory performance for classes like agree and disagree. To address these challenges, our proposed approach integrates CNN and LSTM layers with dimensionality reduction techniques, resulting in superior accuracy of 97.8%.

lll. Methodology:

3.1. Study Design

This systematic literature review follows a well-defined schema commonly employed in the literature [37], consisting of several key components:

Research Questions: Formulating clear research questions to guide the review process.
Search Strategy: Defining the approach for collecting relevant data.
Article Selection: Establishing criteria for selecting studies that meet the objectives of the review.
Distribution of Studies: Analyzing the chronological distribution of selected articles.
Quality Assessment: Evaluating the quality of the studies included in the review.
Data Extraction: Extracting relevant information from the selected articles to address the research questions.

3.2. Research Questions

The research questions (RQs) drive the focus of the study and are designed to explore various aspects of Content-Based Fake News Detection (CBFND). The RQs are as follows:

RQ1: What machine learning models are used in CBFND?
RQ2: What features are used in CBFND machine learning models?
RQ3: What datasets are used in the literature?
RQ4: What are the best-performing algorithms and features used in literature?

3.3. Search Strategy

A search strategy was devised to identify relevant studies for answering the research questions. The search string "(Fake news detection) AND (Machine Learning OR deep learning)" was formulated, utilizing Boolean operators and including synonyms and acronyms to cover a broad range of related literature.

3.4. Article Selection

Four major database repositories were considered for study selection. The search string was applied, resulting in the identification of 145 articles. Exclusion criteria were applied, and duplicates were removed to ensure the integrity of the dataset.

3.5. Quality Assessment

The quality of the selected articles was assessed using predefined criteria inspired by previous literature [38]. Each study was evaluated based on the criteria listed in Table 5. Articles that met five or more criteria were retained for further analysis.

Model Generation Approach

The methodology for generating the fake news classification model involves a fusion of global and local text semantics through three distinct sub-layers:

BERT Processing Layer:

This layer reads the preprocessed data and feeds it into a BERT pre-trained model tuned for embeddings.
BERT processes the text to generate global semantics by analyzing the relationships among current, previous, and upcoming words.
The output from BERT is passed through dense and dropout layers to produce refined representations.
CNN Processing Layer:

The preprocessed text data is inputted into this layer, which converts it into GloVe embeddings.
These embeddings are then passed through three parallel CNN layers, each with kernels of sizes two, three, and four, respectively.
The CNN layers extract local text semantics using n-gram features.
The outputs from the CNN layers are further processed through multiple dense and dropout layers to refine the representations.
Dense Net Processing Layer:

This layer integrates the outputs from the CNN and BERT processing layers, combining both global and local text semantics.
The merged representations are passed through dense and dropout layers for further refinement.
Finally, the model outputs the classification of the news text, labeling it as either real or fake based on the learned features.

lV. Conclusion

In this era of information abundance and digital interconnectedness, the proliferation of fake news presents a formidable challenge to society, impacting discourse, decision-making, and trust in media. Addressing this challenge requires innovative approaches that leverage cutting-edge technologies. In this context, our research introduces a novel benchmark model, WELFake, designed for the detection of fake news. Through meticulous curation of a comprehensive dataset comprising over 72,000 articles from diverse sources, we aimed to ensure unbiased classification outcomes and robust model evaluation.

The WELFake model operates in two distinct phases, each contributing to its overall efficacy in detecting fake news. In the first phase, preprocessing and validation are conducted using linguistic features, laying the foundation for accurate classification. These linguistic features, meticulously analyzed and curated from state-of-the-art works, play a crucial role in discerning subtle patterns indicative of fake news content. By selecting 20 significant linguistic features from over 80 candidates, we minimized computational complexity while maximizing classifier accuracy.

In the second phase, we integrated linguistic feature sets with word embedding (WE) techniques, harnessing the power of both global and local text semantics. The incorporation of WE methods, such as TF-IDF and count vectorization (CV), further enhanced the model's ability to capture nuanced linguistic nuances and contextual information. Through a voting classification scheme, we synthesized the outputs from multiple machine learning models, ensuring robustness and reliability in classification outcomes.

Experimental results on the WELFake dataset demonstrate the remarkable efficacy of our model, achieving an impressive classification accuracy of 96.73%. This represents a significant advancement over existing methodologies, surpassing the accuracy of BERT by 1.31% and CNN models by 4.25%. Moreover, comparative analysis with predictive-based approaches utilizing Word2vec embeddings showcased an improvement of up to 1.73%, underscoring the superiority of our frequency-based analysis of writing patterns.

The findings from our research have far-reaching implications for combating the dissemination of fake news in real-time social media environments. By providing a robust and accurate tool for identifying fake news content, the WELFake model empowers users, media outlets, and policymakers to make more informed decisions and mitigate the negative impacts of misinformation.

Furthermore, our analysis of different machine learning models revealed that SVM produced the most accurate results, highlighting the importance of model selection in achieving optimal performance. Additionally, the WELFake model achieved an overall accuracy of 93%, further reinforcing its reliability and effectiveness in real-world scenarios.

Looking ahead, our research opens up new avenues for further exploration and improvement. Future extensions of our work could incorporate additional factors such as knowledge graphs and user credibility to enhance the verification process further. Additionally, ongoing advancements in deep learning techniques, including XLNet, ALBERT, and LSTM with BERT, hold promise for further improving fake news detection performance.

In conclusion, our study represents a significant contribution to the field of fake news detection, offering a robust and effective solution in the form of the WELFake model. By combining sophisticated linguistic analysis with state-of-the-art machine learning techniques, we have demonstrated the potential to combat the dissemination of fake news and uphold the integrity of information in the digital age.

REFERENCES:

[1] M. H. Goldani, R. Safabakhsh, and S. Momtazi, ‘‘Convolutional neural network with margin loss for fake news detection,’’ Inf. Process. Manage., vol. 58, no. 1, Jan. 2021, Art. no. 102418.

[2] A. Bondielli and F. Marcelloni, ‘‘A survey on fake news and rumour detection techniques,’’ Inf. Sci., vol. 497, pp. 38–55, Sep. 2019.

[3] K. Rapoza, ‘‘Can ’fake news’ impact the stock market,’’ Tech. Rep., 2017. [Online]. Available: https://www.forbes.com/sites/kenrapoza/2017/02/26/ can-fake-news-impact-the-stock-market/?sh=187125e22fac

[4] H. Allcott and M. Gentzkow, ‘‘Social media and fake news in the 2016 election,’’ J. Econ. Perspect., vol. 31, no. 2, pp. 211–236, 2017.

[5] L. Bozarth and C. Budak, ‘‘Toward a better performance evaluation framework for fake news classification,’’ in Proc. Int. AAAI Conf. Web Social Media, vol. 14, 2020, pp. 60–71.

[6] Y. Liu and Y.-F. B. Wu, ‘‘Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks,’’ in Proc. 32nd AAAI Conf. Artif. Intell., Apr. 2018, pp. 354–361.

[7] F. Monti, F. Frasca, D. Eynard, D. Mannion, and M. M. Bronstein, ‘‘Fake news detection on social media using geometric deep learning,’’ 2019, arXiv:1902.06673. [Online]. Available: http://arxiv.org/abs/1902.06673

[8] S. Girgis, E. Amer, and M. Gadallah, ‘‘Deep learning algorithms for detecting fake news in online text,’’ in Proc. 13th Int. Conf. Comput. Eng. Syst. (ICCES), Dec. 2018, pp. 93–97.

[9] S. Ghosh and C. Shah, ‘‘Towards automatic fake news classification,’’ Proc. Assoc. Inf. Sci. Technol., vol. 55, no. 1, pp. 805–807, Jan. 2018.

[10] R. Yamashita, M. Nishio, R. K. G. Do, and K. Togashi, ‘‘Convolutional neural networks: An overview and application in radiology,’’ Insights Imag., vol. 9, pp. 611–629, Aug. 2018.

[11] Y. Yu, C. Wang, X. Gu, and J. Li, ‘‘A novel deep learning-based method for damage identification of smart building structures,’’ Struct. Health Monitor., vol. 18, no. 1, pp. 143–163, Jan. 2019.

[12] H. Li, ‘‘Deep learning for natural language processing: Advantages and challenges,’’ Nat. Sci. Rev., vol. 5, no. 1, pp. 24–26, Jan. 2018.

[13] J. Ma, W. Gao, P. Mitra, S. Kwon, B. J. Jansen, K.-F. Wong, and M. Cha, ‘‘Detecting rumors from microblogs with recurrent neural networks,’’ Tech. Rep., 2016. [Online]. Available: https://ink.library.smu.edu. sg/sis_research/4630/

[14] N. Ruchansky, S. Seo, and Y. Liu, ‘‘CSI: A hybrid deep model for fake news detection,’’ in Proc. ACM Conf. Inf. Knowl. Manage., Nov. 2017, pp. 797–806.

[15] K. Shu, D. Mahudeswaran, S. Wang, D. Lee, and H. Liu, ‘‘FakeNewsNet: A data repository with news content, social context and spatialtemporal information for studying fake news on social media,’’ 2018, arXiv:1809.01286. [Online]. Available: http://arxiv.org/abs/1809.01286

[16] X. Zhang, A.A. Ghorbani, An overview of online fake news: Characterization, detection, and discussion, Inform. Process. Manage. 57 (2020) .

[17] J.A. Nasir, O.S. Khan, I. Varlamis, Fake news detection: A hybrid cnn-rnn based deep learning approach, Int. J. Inform. Manage. Data Insights 1 (2021) .

[18] A. Abedalla, A. Al-Sadi, M. Abdullah, A closer look at fake news detection: A deep learning perspective, in: Proceedings of the 2019 3rd International Conference on Advances in Artificial Intelligence, 2019, pp. 24–28.

[19] S.A. Alameri, M. Mohd, Comparison of fake news detection using machine learning and deep learning techniques, 3rd International Cyber Resilience Conference (CRC), IEEE 2021 (2021) 1–6.

[20] N. Aslam, I. Ullah Khan, F.S. Alotaibi, L.A. Aldaej, A.K. Aldubaikil, Fake detect: A deep learning ensemble model for fake news detection, complexity 2021 (2021)

[21] H. Ahmed, I. Traore, and S. Saad, “Detection of online fake news using N-gram analysis and machine learning techniques,” in Intelligent, Secure, and Dependable Systems in Distributed and Cloud Environments, vol. 10618. Cham, Switzerland: Springer, Oct. 2017, pp. 127–138.

[22] K. Shu, S. Wang, and H. Liu, “Exploiting Tri-relationship for fake news detection,” Dec. 2018, arXiv:1712.07709v1. [Online]. Available: https://arxiv.org/abs/1712.07709v1.

[23]G. Gravanis, A. Vakali, K. Diamantaras, and P. Karadais, “Behind the cues: A benchmarking study for fake news detection,” Expert Syst. Appl.,vol. 128, pp. 201–213, Aug. 2019.

[24]L. Borges, B. Martins, and P. Calado, ‘‘Combining similarity features and deep representation learning for stance detection in the context of checking fake news,’’ J. Data Inf. Qual., vol. 11, no. 3, pp. 1–26, Jul. 2019, doi: 10.1145/3287763.

[25]D. P. Rao, (2017). Exploring How Artificial Intelligence Technologies Could be Leveraged to Combat Fake News. Accessed: Oct. 29, 2019.[Online]. Available: http://www.fakenewschallenge.org/

[26]K. Popat, S. Mukherjee, J. Strötgen, and G. Weikum, ‘‘Where the truth lies: Explaining the credibility of emerging claims on the Web and social media,’’ in Proc. 26th Int. Conf. World Wide Web Companion, Apr. 2017, pp. 1003–1012.

[27]C. Stab, T. Miller, and I. Gurevych, ‘‘Cross-topic argument mining from heterogeneous sources using attention-based neural networks,’’ 2018, arXiv:1802.05758. [Online]. Available: http://arxiv.org/abs/1802.05758

[28] I. Augenstein, T. Rocktäschel, A. Vlachos, and K. Bontcheva, ‘‘Stance detection with bidirectional

conditional encoding,'' in Proc. 2016 Conf. Empirical Methods Natural Lang. Process., Austin, TX, USA, Nov. 2016, pp. 876–885. [Online]. Available: https://www.aclweb.org/anthology/D16-1084

[29]W. Ferreira and A. Vlachos, ''Emergent: a novel data-set for stance classification,'' in Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Hum. Lang. Technol., San Diego, CA, USA, Jun. 2016, pp. 1163–1168. [Online]. Available: https://www.aclweb.org/anthology/N16-1138

[30] Q. Zeng, ''Neural stance detectors for fake news challenge,'' Stanford Univ., Stanford, CA, USA, Tech. Rep., 2017.

[31]S. Mohammad, S. Kiritchenko, P. Sobhani, X. Zhu, and C. Cherry, ''SemEval-2016 task 6: Detecting stance in tweets,'' in Proc. 10th Int. Workshop Semantic Eval. (SemEval-2016), San Diego, CA, USA, Jun. 2016, pp. 31–41. [Online]. Available: https://www.aclweb.org/anthology/S16-1003

[32]C. Stab, T. Miller, and I. Gurevych, ''Cross-topic argument mining from heterogeneous sources using attention-based neural networks,'' 2018, arXiv:1802.05758. [Online]. Available: http://arxiv.org/abs/1802.05758

[33]M. A. Walker, P. Anand, R. Abbott, and R. Grant, ''Stance classification using dialogic properties of persuasion,'' in Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Hum. Lang. Technol., Stroudsburg, PA, USA, 2012, pp. 592–596. [Online]. Available: http://dl.acm.org/citation.cfm?id=2382029.2382124

[34] S. Somasundaran and J. Wiebe, ''Recognizing stances in ideological online debates,'' in Proc. NAACL HLT Workshop Comput. Approaches Anal. Gener. Emotion Text, Los Angeles, CA, USA, Jun. 2010, pp. 116–124. [Online]. Available: https://www.aclweb.org/anthology/W10-0214

[35] G. Bhatt, A. Sharma, S. Sharma, A. Nagpal, B. Raman, and A. Mittal, ''Combining neural, statistical and external features for fake news stance identification,'' in Proc. Companion Web Conf. Web Conf. (WWW), Geneva, Switzerland, 2018, p. 1353, doi: 10.1145/3184558.3191577.

[36]C. Dulhanty, J. L. Deglint, I. Ben Daya, and A. Wong, ''Taking a stance on fake news: Towards automatic disinformation assessment via deep bidirectional transformer language models for stance detection,'' 2019, arXiv:1911.11951. [Online]. Available: https://arxiv.org/abs/1911.11951

[37] J. Biolchini, P.G. Mian, A.C.C. Natali, G.H. Travassos, Systematic review in software engineering, System engineering and computer science department COPPE/UFRJ, Technical Report ES 679 (2005) 45.

[38] M. Galster, D. Weyns, D. Tofan, B. Michalik, P. Avgeriou, Variability in software systems–a systematic literature review, IEEE Trans. Software Eng. 40 (2013) 282–306.