Automatic detection of Fake News

Pontus Nordberg[a], Joakim Kävrestad[a] and Marcus Nohlberg[a]

[a]University of Skövde, Högskolevägen 1, 541 28 Skövde, Sweden

Abstract

Following the American presidential election in 2016, the terms ”fake news” was popularized and has since been a common term in the public vocabulary. While quite recently popularized, fake news is a phenomenon that is as old as news itself and is most commonly defined as purposeful disinformation used to untrue information or skewed reporting intended to push a certain narrative. In recent years, fake news has seen frequently in attempts to influence elections or by organized crime organizations in various efforts to make money, not least drawing from the ongoing CoVid-19 pandemic. We argue that the phenomenon must be researched from technical as well as from social aspects, since it involved using technical tools to spread information targeted humans. In this paper, we identify key methods for automatic fake news detection in order to lay the foundation for end-user support system designed to help users identify and avoid fake news.

Keywords

fake news, machine learning, classification, automatic, disinformation

1. Introduction

In the lead-up year before the 2016 American presidential election the term "fake news” was popularized and has since become a common term in the public vocabulary. Views have differed in regard to exactly what content should or should not be included under the "fake news”-umbrella. However, the most commonly used definition is that fake news consists of deliberate disinformation, used to purposefully spread untrue information, or as skewed reporting of real events to push a certain narrative [1].

While the concept of fake news is as old as news itself, the advent of social media and mass-information on the Internet has led to fake news taking on a new form compared to its previous iterations [2]. Concerns have been raised especially in the context of the effect which fake news can have on elections [3]. Another effect of fake news is destabilization where foreign states spread fake news in other states with the purpose of destabilising democracy, this has been seen in the ongoing CoVid-19 pandemic [4].

We argue that Fake News is a socio-technical phenomenon that aims to exploit human behaviour but relies on technical infrastructures and services in order to spread. The same perspective is used by [5] who discuss reasons for why users share fake news and provide a social technical model of media effects. We argue that researching Fake News from socio-technical perspective allow for a more holistic view than just looking at the technical or social aspects, and is necessary to fully understand the domain.

Much of the research that has been done as a result of fake news becoming an increasingly relevant issue has been aimed towards finding ways to combat it. One way to combat Fake News could be to leverage "in-the-moment", such as ContextBased MicroTraining [6], that identifies information that

6th International Workshop on Socio-Technical Perspective in IS development (STPIS'20), June 08–09, 2020, Online

EMAIL: joakim.kavrestad@his.se (J. Kävrestad); marcus.nohlberg@his.se (M. Nohlberg)

ORCID: 0000-0003-2084-9119 (J. Kävrestad); 0000-0001-5962-9995 (M. Nohlberg)

could be fake news and provides the user with information the user needs in order to determine if the news she is reading is legitimate or fake.

This study seeks to support such an endeavour by identifying and evaluating existing methods for automatic fake news detection. The results will provide insight into how Fake News can be automatically detected and be used in future development of user support tools. The study will also provide an overview of existing research in the area.

2. Methodology

The study was conducted by means of a structured literature review, as described by [7]. The search term *Automatic Fake News Detection* was used in the following three databases:

- IEEExplore
- SpringerLink
- ACM Digital Library

Google Scholar was used for complementary searches, as suggested by [8]. Following the searches, the resulting papers were evaluated against the following inclusion criteria:

- Papers should be published 2015 or later
- Papers should be written in English
- Papers should present or evaluate automatic detection methods for Fake News
- Papers should be peer-reviewed

The limitation in time was chosen to get results from the most recent research and research following the exploding interest in the term *Fake News* following the presidential election in the US. The rest of the exclusion criteria were chosen to ensure that only high-quality papers related to the topic of this study were included. The search process and application of exclusion criteria resulted in *47* papers that were included in the study.

The included papers were analyzed using thematic coding as described by [9]. They were read to identify what type of automatic detection they presented, and the presented detection types were abstracted to distinct themes that reflected methods for automatic fake news detection.

3. Results

After coding all accepted papers, four themes emerged as follows:

- *Text Classification* – Papers mainly focusing on analysis and classification based on the content of the article (25 papers)
- *Network* – Papers mainly focusing on how the fake news is spread (10 papers)
- *Human-Machine Hybrid* – Supplementing automatic detection with manual human input (four papers)
- *Reviews & Evaluations* – reviews of the current automatic detection methods (eight papers)

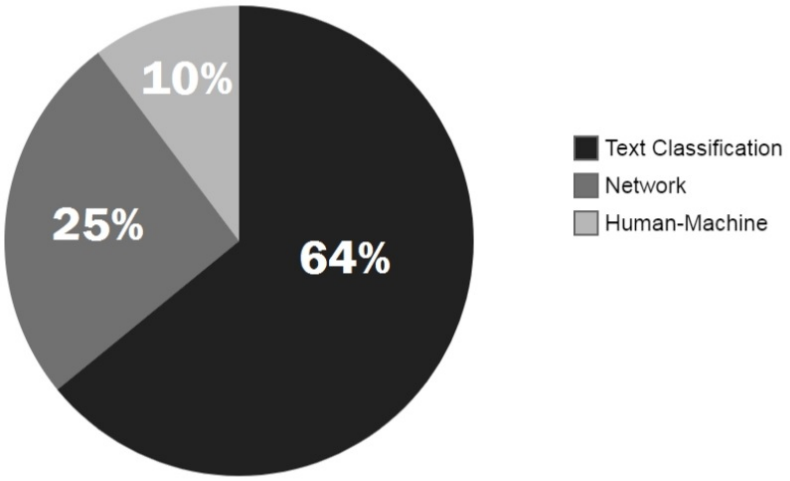

Figure 1: Distribution of papers in the identified themes

Theme	Papers
Text Classification (25) – Papers mainly focusing on analysis and classification based on the content of the article	[10, 11, 12, 13, 2, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 3, 24, 25, 26, 27, 28, 29, 30, 31, 32]
Network Analysis (10) – Papers mainly focusing on how the fake news is spread	[33, 34, 35, 36, 37, 38, 39, 40, 41, 42]
Human-Machine Hybrid (4) – Supplementing automatic detection with manual human input	[43, 44, 45, 46]
Reviews and evaluations (8) – reviews of current automatic detection methods	[47, 48, 49, 50, 51, 52, 53, 54]

Table 1
List of papers included in this study

The first three themes are identified methods used for automatic fake news detection, and the last theme include papers evaluating those methods. The most prominent method is to base fake news detection on *text classification*, a method that has shown good results in several studies. However, *text classification* is challenging in some cases, for instance, when little text is available. As such, the conclusion of this study is that the best approach for automatic fake news detection is a combination of text classification and a *Network* approach. The distribution of identified papers in the three first themes is visualized in Figure 1, followed by Table 1, listing the papers included in the study. The rest of this section will discuss the identified methods in more dept.

3.1. Text Classification

As can be seen in Figure 1, above, Text classification is the most popular approach of automated fake news detection and the majority of the collected papers propose solutions using such methods. A large number of researches hold that automatic detection of fake news can achieve promising results when the problem is addressed by machine learning and deep learning classification, but that it still needs work in certain areas in order to become fully reliable. Researches have concluded that many classifier tools can be effective for detecting fake news, achieving high accuracy.

Several authors show the effectiveness of classifiers and models developed with these. [20] examine several machine learning techniques in the form of different classifiers which all display a high level of accuracy when used on fake news datasets, the most effective being the Linear SVM (Support Vector Machine) classifier. [3] also utilized SVM based algorithms with their One-Class Classification approach, which was concluded to have the potential to not only distinguish fake news from real news, but also from opinion-pieces and propaganda. [31] showed good results with the GWO (Grey Wolf Optimization) algorithm. [16] showed how various machine learning classifiers were integrated into a single multi-level model in which they could increase performance and results by working together, helping to offset each other's weaknesses. Classification relying on lexical rules, syntax, semantics, and similar factors achieved promising results, with the models being able to achieve results equal to that of the human ability to detect fake news.

The authors of the papers in this category discuss many different aspects of machine learning classification and how it can be used with the methods they present. For example, eye-catching headlines are a common feature in fake news and is therefore often a target for analysis. Text classification tools are presented as being useful for detecting fake news by analysing the relation between the article headline and the body. It is also argued, however, that such a method fails to exploit the article structure, and enhanced text classification is presented. It gives differential importance to sections of an article using the context and coherence of words, sentences, and their relation to the headline to better detect fake news based on headlines. Other authors also follow this approach. Using similar types of classifiers together with "textual entailment" features and "stance detection". They also analyse the article body and its headline's relationship and if the latter is related or unrelated to the corresponding claims by examining how much they "agree" or "disagree" with each other. This data can also be combined with what other news agencies are saying about the same report.

[23] presented the similarity shared by between fake news articles and spam messages, regarding properties such as syntactic similarities, grammatical mistakes, emotional tone, and trying to affect the reader's opinion. Utilizing simple classifiers often used for spam filtering against fake news, they achieved promising results relative to the simplicity of the model used. [25, 30] focus on the emotional tone in particular, common in fake news's attempts to trigger intense emotions to achieve further spread online. They present a detection model which incorporates extracted emotional signals, such as the use of emotional words in a claim, the "emotional intensity" expressed by a claim, and a neural network that predicts the emotional intensity that can be triggered in a user. [27] use a Convolutional Neural Network model which analyses the semantic and syntactic information from articles in addition to text analysis. This was used to calculate the weight of "sensitive" words, and showed high results.

A problem mentioned by [18] is the inability of some models to properly filter fake news in Portugal due to the scarcity of labelled datasets in Portuguese, which leads to the necessity of having to develop their own linguistic reference. [32] developed a dataset in German for the same reason. [29] aim at easily developing a large enough dataset to be used in the creation of an automatic classification model, which they regard as the true challenge as opposed to the, technically speaking, rather straight-

forwards classification problem of fake news detection itself. Their solution is to automatically collect hundreds of thousands of tweets into a dataset and label their source as trustworthy or untrustworthy, training a classifier with that dataset, and then using that classifier to classify other tweets as true or false. Despite all tweets from an untrustworthy source not being fake and vice versa, the resultant datasets could still achieve rather high performance, making dataset acquisition a simple process if a certain amount of label noise is accepted.

[13] proposes a method which, instead of being based on traditional deep learning, is based on "active learning", which can efficiently learn from smaller samples in order to improve performance in weakly-supervised learning environments. Making it more suited for practical applications such as when dealing with 140 character length twitter posts.

[2] utilized machine learning classifiers to identify fake news by examining visual cues based on things such as image content, brightness and "mood" combined with the content of the text; how it is written, usage of positive or negative words, as well as with behavioural cues in regards to how news are shared online, to automatically classify it as fake news or not.

[24] looked at satirical fake news, which they claim can become harmful in the existing atmosphere of fake news, despite not being true "fake news" due to their deliberately humorous angle. Reasons for that is the fact that many users only read the headline of articles they come across, combined with shared/liked articles in social media such as Facebook also only showing the headlines in the newsfeed. They look at how humorous satirical news is constructed and uses a machine learning approach that aims to automatically detect them.

3.2. Network Analysis

As some of the papers in the above category start to touch on, the surrounding data outside of the fake news article itself can also be an essential part of automatic fake news detection. This is the so-called "Network" approach, which can be an effective tool and has some advantages in areas where the text classification approach can run into issues. Several researchers also include features specifically related to social media, the largest area of distribution of fake news, in the detection models presented within this area. How fake news is propagated in such environments, as well as what role users play in this, is particularly in focus.

The challenge of detecting fake news with limited information is a recurring theme in the literature under this category. Machine learning approaches which detect fake news based on text analysis have a limitation in the form of not being able to detect fake news early, when the information required for verification or debunking is unavailable due to the early stage of news propagation. This limits the detection accuracy, and a model is proposed which is based on a Convolutional Network that detects the propagation paths of the news based on the users who spread it. Detection based on global and local user variations along this path can then help identify fake news with high accuracy much earlier than other approaches. Text analysis is also too ineffective when news pieces are too short, but this can also be analyzed using convolutional networks. They can create knowledge graphs which can detect the veracity of the relation between two entities. Such models would generate background information knowledge graphs which can then be used to automatically build entity relations to detect truth or falsehood. The use of structured knowledge networks is suggested to be able to gather the background information. Despite a low amount of information being used to create knowledge graphs, they were proved to be able to provide decent results.

[33] also used knowledge networks to detect deception. False statements being presented as real can be extracted and examined alongside statements which can be found in structured knowledge networks online. A statement would then have its subject and predicate divided into nodes which,

based on their "proximity" in terms of narrow or wide relation to each other, would increase the likelihood that the type of pairing of the subject-predicate in the statement is true or false. In this way, fact-checking can become a simple computation to detect the shortest "path" between nodes in a statement. Such an approach would however require a relevant pre-existing knowledge base. In social media, metadata and included hyperlinks can be used to establish veracity by linking the most important words to other words within the same network.

[50] propose a network-based approach which studies the patterns of how the news are being spread, who spreads them, and what the relation between these spreaders are. Such patterns are then used at various network levels to detect fake news. The authors show how fake news spreads farther than real news, and has more and stronger spreader engagement, which is denser for networks of similar interests and behaviour. A downside is that the news needs to be propagated on social media before it can be detected, although here as well only a small amount of network information is needed.

Another approach which further focuses on the user engagement uses a combination of both the text-analysis in addition to a social-based method which takes into account the amount of times content is shared or "liked", as well as by who. It has been shown that content of Facebook can be classified as fake news with very high accuracy based on the users who "like" them, a result of the tendency of users to aggregate into "bubbles" which selectively share content that fits their narrative of things. A downside to this is of course that the content needs to have a certain amount of "social interaction" in order to produce worthwhile results..

[38] also propose a model that looks at features like the group behaviour of users who spread fake news, and from who and where they spread it from. This would be done through a combined text analysis and network analysis model which combines the three fake news characteristics of: the article text, the user response it receives, and the user source promoting it. A combined approach like this is argued to increase accuracy and generality, managing to achieve this as well as a meaningful representation of users and articles.

Combination approaches were also proposed by [41], who utilized machine learning classifiers and demonstrated the importance of combining statements with the overall "credibility patterns" of politicians. Similarly, [40] created a diffusive network model that relies on explicit and latent features extracted from a text and achieved very good performance in not only identifying fake news article, but also striking at the origin of such news by likewise identifying the news creators and types of news subjects, where a record of credibility can be seen and assessed. The latter two categories being considered more important than the news itself in the contribution to eradicate fake news spread on social media.

3.3. Human-Machine Hybrid

Some of the collected literature proposed a human-machine hybrid method. Such methods utilize the machine-based automatic detection models mentioned above, but also rely on manual input from humans in various ways to help improve accuracy and performance. A hybrid method is hoped to bring the best of both worlds. For example, [44] examine how human, machine and human-machine-based approaches compare and find the hybrid approach to be most effective.

[46] state that automatic fake news detection is made difficult due to the language used in such articles being purposefully hard to distinguish from real stories, and differentiating satirical news from fake news is even more complicated. Classical fake news detection, such as linguistic approaches, does not work well with satirical news about politics due to the advanced language often used. Their solution is using a machine learning-human crowd hybrid model, where human input is used depending on the confidence of an initial machine classification. They achieve better results than a pure

machine approach, and a combined method also helps alleviate the poor scalability in regards to cost and latency that a purely crowd-sourced approach would suffer from. They suggest such a system could be applied to social media.

[45] propose a machine learning model that crosschecks content input with articles from "Fact DB", described as a collection of true articles, built and updated by human's direct judgment for detecting fake news by collecting facts. The input would then be compared with content from Fact DB using sentences matching. An article abstraction process and an entity matching set (alongside the aforementioned sentence matching) are used to improve performance and accuracy when dealing with longer sentences or unlearned relations between words, respectively. A final answer regarding the factuality of the input is produced based on the results of the two matching steps.

[43] propose a very ambitious solution in the form of a fact-checking infrastructure which would be designed as an additional layer on top of the World Wide Web, available to every user, ideally integrated into all web browsers to the same degree as the URL-field or a bookmark function, and able to handle millions of users, arbitrary types of content, many different languages and massive amounts of data. The infrastructure would provide web-based tools that users could use to process any content on a page in order to get more information. A user would target any content, ranging from a single comment to an entire article spread over several pages, and receive additional/alternate viewpoints, independent assessments, or indications if the content is dangerous/abusive/factual. Automatic machine learning methods combined with human intelligence and feedback are considered necessary to fulfil these tasks, with a fully automatic or hybrid approach being used for different areas of analysis. The infrastructure should be decentralised to avoid being vulnerable to misuse. Any organisation/company/research centre should be able to operate and offer services to the infrastructure and users should be able to configure a personal set of these (fully transparent) tools and services, fully combinable with each other using standardised in/output formats, in order to aggregate a value regarding, for example, political bias. Decentralised web annotation repositories, aggregated from both human and machine, would send their annotations to the tools which provide information to the end-user.

3.4. Reviews and evaluation

This category includes literature which does not propose a specific solution or method, but rather looks at several methods and compares and evaluates them.

[48] performed a comparative study on several types of text classification and language processing machine learning algorithms. They concluded that the XGB (Gradient Boosting Algorithm) classifier, which implements machine learning algorithms, was the most effective classifier of the ones they tested.

[47] stresses the limitation of existing datasets, the quality of which various detection models depend on, due to the reliance on human annotations and instead propose a method for automation of the data set creation process. They create an automatic system which is based on fake-news stories classified by face-checking websites. These types of datasets are important for all kinds of detection models.

[54] examines the proposed approach of tracing a text to its source based on the writing style of the article and determining the trustworthiness of that source. The reason for this approach is the fact that much of fake news is more or less auto-generated, and this can be identified with text analysis and can be indicated by the type of source itself. However, since true news can be auto-generated in a similar process, or that fake news can consist of corrupted real news, the authors reveal that this approach has a big weakness. This highlights the importance of assessing the veracity of the text, and

not just its style and source.

[49] evaluate different tools using the same dataset for all of them. The tested tools are SurfSafe; a browser extension which compares images from news with a database; Fake News Detector AI, which uses a neural network; TrustedNews, which can detect bias, clickbait, satire, etc.; Fake News Detector, that is open-source uses feedback from users of the tool; Fake News Guard, a browser extension which verifies any page visited by the user or link displayed in Facebook using linguistic and network analysis plus AI; Decodex, a tool which can label information as “info”, “satire”, or “no info”. Many of the classification attempts failed, producing only responses like “error” or “no result”. There were however almost no cases of different tools providing contradictory classifications the same piece of information (i.e. true and false at the same time). Results were better when using English text and sources.

There were several papers that looked at fake news detection itself, such as its challenges, research directions, characterisation and comparison versus other concepts such as satire or rumours. They talk about the different detection techniques and mechanisms which exist. The information brought up in these areas have however already been covered in the chapters above.

4. Conclusions and Discussion

The aim of this paper was to summarise existing methods for automatic fake new detection. The study find that there is a lot of research in the domain of fake news detection and while there are several methods available, they all have benefits and drawbacks. The examined literature reveals attributes that can commonly be out on fake news and used for identification, including:

- Fake news are often emotionally charged
- Fake news often have eye-catching headlines that can be compared to click bait
- Fake news spread patterns appear to be pre-determinable

While these traits apply to fake news, they can also apply to there content types making it difficult to distinguish fake news from, for instance, satiric content. In conclusion, it seems feasible to base fake news detection on not only text content but also on spread patterns, as suggested by for instance [33, 39, 41].

This paper summarizes methods for automatic Fake news detection discussed in recent scientific literature. As such, it contributes to the scientific community as a snapshot of the current research landscape. It can also benefit practitioners seeking to develop automatic classification algorithms for use in Fake news prevention.

Under underlying question, raised by the topic itself, is how Fake news detection algorithms should be implemented in practice. It would seem necessary to apply a permissive or restrictive approach where a permissive approach would limit the amount of true news that are classified as fake and a restrictive approach would perhaps correctly identify more fake news but also classify legitimate content as fake. Another important implementation factor is how the algorithms will be perceived by the users. A Fake news detection software will undeniably be a socio-technical system designed to assist Internet users. As such, studies concerning the social aspect of fake news detection is much needed as future work.

While the actual implementation of Fake News detection is beyond the scope of this paper, the topic raises an delicate ethical dilemma. Any automatic Fake News detection will essentially work as a filter and one must consider who that should decide what is classified as legitimate and what that

is classified as fake or fraudulent [55]. As such, an important direction for future work is the domain of ethics and content filtering. A common framework on this topic is much needed by researchers as well as practitioners. We foresee a key question to be concerning who decides that is fake and what is legitimate. In light of the US presidential election campaign of 2020, it is evident that different stakeholders have different perspectives on this matter.

References

[1] A. Gelfert, Fake news: A definition, Informal Logic 38 (2018) 84–117.

[2] C. Janze, M. Risius, Automatic detection of fake news on social media platforms., in: PACIS, 2017, p. 261.

[3] P. Faustini, T. F. Covões, Fake news detection using one-class classification, in: 2019 8th Brazilian Conference on Intelligent Systems (BRACIS), IEEE, 2019, pp. 592–597.

[4] S. Boberg, T. Quandt, T. Schatto-Eckrodt, L. Frischlich, Pandemic populism: Facebook pages of alternative news media and the corona crisis–a computational content analysis, arXiv preprint arXiv:2004.02566 (2020).

[5] A. E. Marwick, Why do people share fake news? a sociotechnical model of media effects, Georgetown Law Technology Review 2 (2018) 474–512.

[6] J. Kävrestad, M. Nohlberg, Using context based micro training to develop oer for the benefit of all, in: Proceedings of the 15th International Symposium on Open Collaboration, 2019, pp. 1–10.

[7] P. Brereton, B. A. Kitchenham, D. Budgen, M. Turner, M. Khalil, Lessons from applying the systematic literature review process within the software engineering domain, Journal of systems and software 80 (2007) 571–583.

[8] B. V. Wee, D. Banister, How to write a literature review paper?, Transport Reviews 36 (2016) 278–288.

[9] R. Anderson, Thematic content analysis (tca), Descriptive presentation of qualitative data (2007).

[10] S. Singhania, N. Fernandez, S. Rao, 3han: A deep neural network for fake news detection, in: International Conference on Neural Information Processing, Springer, 2017, pp. 572–581.

[11] 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.

[12] T. Saikh, A. Anand, A. Ekbal, P. Bhattacharyya, A novel approach towards fake news detection: deep learning augmented with textual entailment features, in: International Conference on Applications of Natural Language to Information Systems, Springer, 2019, pp. 345–358.

[13] S. D. Bhattacharjee, A. Talukder, B. V. Balantrapu, Active learning based news veracity detection with feature weighting and deep-shallow fusion, in: 2017 IEEE International Conference on Big Data (Big Data), IEEE, 2017, pp. 556–565.

[14] V. Pérez-Rosas, B. Kleinberg, A. Lefevre, R. Mihalcea, Automatic detection of fake news, arXiv preprint arXiv:1708.07104 (2017).

[15] E. Qawasmeh, M. Tawalbeh, M. Abdullah, Automatic identification of fake news using deep learning, in: 2019 Sixth International Conference on Social Networks Analysis, Management and Security (SNAMS), IEEE, 2019, pp. 383–388.

[16] S. Kaur, P. Kumar, P. Kumaraguru, Automating fake news detection system using multi-level voting model, Soft Computing (2019) 1–21.

[17] T. Traylor, J. Straub, N. Snell, et al., Classifying fake news articles using natural language processing to identify in-article attribution as a supervised learning estimator, in: 2019 IEEE 13th

International Conference on Semantic Computing (ICSC), IEEE, 2019, pp. 445–449.
[18] R. A. Monteiro, R. L. Santos, T. A. Pardo, T. A. De Almeida, E. E. Ruiz, O. A. Vale, Contributions to the study of fake news in portuguese: New corpus and automatic detection results, in: International Conference on Computational Processing of the Portuguese Language, Springer, 2018, pp. 324–334.
[19] S. Girgis, E. Amer, M. Gadallah, Deep learning algorithms for detecting fake news in online text, in: 2018 13th International Conference on Computer Engineering and Systems (ICCES), IEEE, 2018, pp. 93–97.
[20] H. Ahmed, I. Traore, S. Saad, Detection of online fake news using n-gram analysis and machine learning techniques, in: International conference on intelligent, secure, and dependable systems in distributed and cloud environments, Springer, 2017, pp. 127–138.
[21] S. I. Manzoor, J. Singla, Fake news detection using machine learning approaches: A systematic review, in: 2019 3rd International Conference on Trends in Electronics and Informatics (ICOEI), IEEE, 2019, pp. 230–234.
[22] L. Waikhom, R. S. Goswami, Fake news detection using machine learning, Available at SSRN 3462938 (2019).
[23] M. Granik, V. Mesyura, Fake news detection using naive bayes classifier, in: 2017 IEEE First Ukraine Conference on Electrical and Computer Engineering (UKRCON), IEEE, 2017, pp. 900–903.
[24] V. L. Rubin, N. Conroy, Y. Chen, S. Cornwell, Fake news or truth? using satirical cues to detect potentially misleading news, in: Proceedings of the second workshop on computational approaches to deception detection, 2016, pp. 7–17.
[25] A. Giachanou, P. Rosso, F. Crestani, Leveraging emotional signals for credibility detection, in: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2019, pp. 877–880.
[26] P. Ksieniewicz, M. Choraś, R. Kozik, M. Woźniak, Machine learning methods for fake news classification, in: International Conference on Intelligent Data Engineering and Automated Learning, Springer, 2019, pp. 332–339.
[27] Q. Li, Q. Hu, Y. Lu, Y. Yang, J. Cheng, Multi-level word features based on cnn for fake news detection in cultural communication, Personal and Ubiquitous Computing (2019) 1–14.
[28] S. Alam, A. Ravshanbekov, Sieving fake news from genuine: A synopsis, arXiv preprint arXiv:1911.08516 (2019).
[29] S. Helmstetter, H. Paulheim, Weakly supervised learning for fake news detection on twitter, in: 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), IEEE, 2018, pp. 274–277.
[30] O. Ajao, D. Bhowmik, S. Zargari, Sentiment aware fake news detection on online social networks, in: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2019, pp. 2507–2511.
[31] F. A. Ozbay, B. Alatas, A novel approach for detection of fake news on social media using metaheuristic optimization algorithms, Elektronika ir Elektrotechnika 25 (2019) 62–67.
[32] I. Vogel, P. Jiang, Fake news detection with the new german dataset "germanfakenc", in: International Conference on Theory and Practice of Digital Libraries, Springer, 2019, pp. 288–295.
[33] N. K. Conroy, V. L. Rubin, Y. Chen, Automatic deception detection: Methods for finding fake news, Proceedings of the Association for Information Science and Technology 52 (2015) 1–4.
[34] J. Z. Pan, S. Pavlova, C. Li, N. Li, Y. Li, J. Liu, Content based fake news detection using knowledge graphs, in: International semantic web conference, Springer, 2018, pp. 669–683.
[35] Y. Liu, Y.-F. B. Wu, 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.
[36] G. Hu, Y. Ding, S. Qi, X. Wang, Q. Liao, Multi-depth graph convolutional networks for fake news detection, in: CCF International Conference on Natural Language Processing and Chinese Computing, Springer, 2019, pp. 698–710.
[37] X. Zhou, R. Zafarani, Network-based fake news detection: A pattern-driven approach, ACM SIGKDD Explorations Newsletter 21 (2019) 48–60.
[38] N. Ruchansky, S. Seo, Y. Liu, Csi: A hybrid deep model for fake news detection, in: Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, 2017, pp. 797–806.
[39] M. L. Della Vedova, E. Tacchini, S. Moret, G. Ballarin, M. DiPierro, L. de Alfaro, Automatic online fake news detection combining content and social signals, in: 2018 22nd Conference of Open Innovations Association (FRUCT), IEEE, 2018, pp. 272–279.
[40] J. Zhang, B. Dong, S. Y. Philip, Deep diffusive neural network based fake news detection from heterogeneous social networks, in: 2019 IEEE International Conference on Big Data (Big Data), IEEE, 2019, pp. 1259–1266.
[41] F. C. Fernández-Reyes, S. Shinde, Evaluating deep neural networks for automatic fake news detection in political domain, in: Ibero-American Conference on Artificial Intelligence, Springer, 2018, pp. 206–216.
[42] T. Hamdi, H. Slimi, I. Bounhas, Y. Slimani, A hybrid approach for fake news detection in twitter based on user features and graph embedding, in: International Conference on Distributed Computing and Internet Technology, Springer, 2020, pp. 266–280.
[43] G. Rehm, An infrastructure for empowering internet users to handle fake news and other online media phenomena, in: International Conference of the German Society for Computational Linguistics and Language Technology, Springer, 2017, pp. 216–231.
[44] E. M. Okoro, B. A. Abara, A. Alan-Ajonye, Z. Isa, A. Umagba, Effects of human and human-machine fake news detection approaches on user detection performance, International Journal of Advanced Research in Computer Science 10 (2019) 26.
[45] K.-H. Kim, C.-S. Jeong, Fake news detection system using article abstraction, in: 2019 16th International Joint Conference on Computer Science and Software Engineering (JCSSE), IEEE, 2019, pp. 209–212.
[46] S. Shabani, M. Sokhn, Hybrid machine-crowd approach for fake news detection, in: 2018 IEEE 4th International Conference on Collaboration and Internet Computing (CIC), IEEE, 2018, pp. 299–306.
[47] D. P. Karidi, H. Nakos, Y. Stavrakas, Automatic ground truth dataset creation for fake news detection in social media, in: International Conference on Intelligent Data Engineering and Automated Learning, Springer, 2019, pp. 424–436.
[48] A. P. S. Bali, M. Fernandes, S. Choubey, M. Goel, Comparative performance of machine learning algorithms for fake news detection, in: International Conference on Advances in Computing and Data Sciences, Springer, 2019, pp. 420–430.
[49] A. Giełczyk, R. Wawrzyniak, M. Choraś, Evaluation of the existing tools for fake news detection, in: IFIP International Conference on Computer Information Systems and Industrial Management, Springer, 2019, pp. 144–151.
[50] X. Zhou, R. Zafarani, K. Shu, H. Liu, Fake news: Fundamental theories, detection strategies and challenges, in: Proceedings of the twelfth ACM international conference on web search and data mining, 2019, pp. 836–837.
[51] Z. I. Mahid, S. Manickam, S. Karuppayah, Fake news on social media: Brief review on detection techniques, in: 2018 Fourth International Conference on Advances in Computing, Communica-

tion & Automation (ICACCA), IEEE, 2018, pp. 1–5.
[52] J. C. Reis, A. Correia, F. Murai, A. Veloso, F. Benevenuto, Supervised learning for fake news detection, IEEE Intelligent Systems 34 (2019) 76–81.
[53] 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.
[54] T. Schuster, R. Schuster, D. J. Shah, R. Barzilay, Are we safe yet? the limitations of distributional features for fake news detection, arXiv preprint arXiv:1908.09805 (2019).
[55] E. Mustafaraj, P. T. Metaxas, The fake news spreading plague: was it preventable?, in: Proceedings of the 2017 ACM on web science conference, 2017, pp. 235–239.