SBC Reviews on Computer Science, 2023, 4:1, doi: 10.5753/reviews.2024.2718

The Use of Syntactic Information in Fake News Detection: A Systematic Review

Matheus José Garcia Fagundes [**University of São Paulo** | ***matheus.fagundes@usp.br***]
Norton Trevisan Roman [**University of São Paulo** | ***norton@usp.br***]
Luciano Antonio Digiampietri [**University of São Paulo** | ***digiampietri@usp.br***]

Escola de Artes, Ciências e Humanidades da Universidade de São Paulo Rua Arlindo Bettio, 1000 - Ermelino Matarazzo, 03828-000, São Paulo - SP - Brasil

Received: 07 June 2023 • **Accepted:** 18 January 2024 • **Published:** 07 March 2024

Abstract Fake news has been a critical problem for society, to the extent that its damaging effects can already be seen in several areas, such as democracy and health. However, as fake news grow in number, manual fact-checking becomes impractical for identifying them, which makes automatic detection a compelling alternative. In this sense, this study gathers multiple solutions for the problem of automatically detecting fake news, through the usage of both lexical and syntactic information. This study consists of a systematic review on fake news detection through linguistic patterns, focusing on the use of syntax to aid in the task. Solving complex problems by capturing linguistic patterns is mostly explored in the Natural Language Processing (NLP) area. In general, the use of shallow syntax representations, such as Parts of speech, only marginally increases the performance of classifiers in this task. However, relying on deeper syntactic representations, such as context-free grammars or syntactic dependency trees, present more promising results.

Keywords: Fake news, Syntax, Syntactic information

1 Introduction

Fake news has been harming society in various aspects, insomuch that they are seen today as one of the greatest threats to democracy, journalism, and freedom of expression [Zhou *et al.*, 2019]. Within this setup, social networks have been a powerful source of fake news dissemination [Shu *et al.*, 2017a], sometimes spreading faster than real news through some social media, such as Twitter for example [Vosoughi *et al.*, 2018].

Fake news is defined as fabricated information that mimics the content of a news media in their format but not in their organizational process or intent [Lazer *et al.*, 2018]. According to Rubin *et al.* [2015], they can be classified into three categories: a) Serious fabrication, b) Large-scale rumor, and c) Satire. The difference between them lies in that, whereas the first two are intended to deceive the reader, the third aims at entertainment and humor.

The exact measure of the destructive impact of fake news is difficult to estimate, but damages to democracy can already be noticed. It is estimated that in the 2016 US election, an average adult saw one or more fake news stories during the election month and more than half remember seeing and believing them [Allcott and Gentzkow, 2017]. As another example, in Italy a quantitative observational study was carried out to measure how much "fake news" and corresponding verified news has circulated during the early Covid-19 pandemic [Moscadelli *et al.*, 2020]. The study found that links containing fake news accounted for 23.1% – almost a quarter – of the total amount of shares.

Detecting fake news, however, becomes increasingly unmanageable and the manual checking process is difficult and quite laborious. Therefore, it becomes necessary to develop an automatic mechanism to assist fact-checkers in mitigating the problem of fake news. Such a mechanism could, for example, be incorporated by digital platforms to improve the trustworthiness of their content.

To this end, existing mechanisms usually approach the problem according to four different strategies [Zhou and Zafarani, 2020]: (I) content-based analysis, which consists of fact-checking; (II) style-based analysis, which checks for patterns of how fake news are written; (III) analysis based on propagation, which checks the pattern of propagation; and (IV) analysis based on the source, which assesses the credibility of the source.

The present study can be included in category (II), which consists of fake news detection through the analysis of the different linguistic patterns used both in fake and real news. This is a strategy that is usually explored through the use of Natural Language Processing (NLP) techniques. In this study, we focus on the subset of techniques that also rely on syntactic information for fake news detection.

On this regard, some existing surveys have already explored the problem of fake news detection through various approaches. Shu *et al.* [2017b], for example, provide a comprehensive review of detecting fake news on social media, mentioning the usage of some syntactic information such as probabilistic context for grammars (PCFG). However, due to its broad scope, it does not delve deeper into syntax or present the results achieved with its use. Similarly, Oshikawa *et al.* [2020] presents a systematic review of task formulations, datasets, and NLP solutions for fake news detection, focusing more on the models and corpora for classification without specifically mentioning the usage of syntactic information.

In contrast, Zhou and Zafarani [2020] reviews and evaluates methods that can detect fake news following the four aforementioned strategies, describing some syntactic methods such as Part of Speech (POS), categorized as shallow syntax, and PCFG, categorized as deep syntax. They also present promising results on the performance of these features. While Zhou and Zafarani [2020] covers some methods that rely on syntactic information, there is still room for further exploration of the impact this kind of information has on fake news detection.

Regarding the above mentioned distinction between deep and shallow syntax, in this work we adopt the nomenclature observed during our research, whereby shallow syntax refers to methods that focus on the word-level only whereas deep syntax is used to categorize entire sentences, being therefore more complex and complete [Zhou and Zafarani, 2020; Zhou *et al.*, 2020]. Examples of shallow syntax is Parts-of-Speech tagging (*e.g.* Zewdu and Yitagesu [2022]), while deep syntax comprises methods such as Probabilistic Context Free Grammar (*e.g.* Klein and Manning [2003]) and Dependency Tree (*e.g.* Moschitti [2006]).

Overall, syntactic features are only slightly explored in the field of fake news detection, in contrast to lexical information. When they are used, however, such as in Zhou and Zafarani [2020], they usually show good results. Related articles suggest that future work should explore more syntactic and semantic features to improve classification performance (*e.g.* [Santos *et al.*, 2020; Pisarevskaya, 2017]. Furthermore, combining different strategies has shown promising results for detecting fake news [Silva *et al.*, 2020].

Considering the opportunity to further assess the value of syntax in the classification of fake news, we carry out a systematic literature review to provide a concise assessment of the usage and effectiveness of syntactic information in automatic fake news detection. The rest of this article is organized as follows. Section 2 presents the method we followed for collecting the articles that comprise our study. Next, in Section 3, we describe the conduction process of the systematic review. Section 4, in turn, shows the review's results, with a deeper discussion being presented in Section 5. Finally, in Section 6 we present our final remarks.

2 Methodology

The methodology of this study involves conducting a systematic literature review.. As described by Kitchenham [2004], a systematic literature review is a means of identifying, evaluating, and interpreting all available research relevant to a particular research question, topic area, or phenomenon of interest. This methodology follows a predefined protocol, allowing for its reproduction, provided all search mechanisms do not change their behavior over time.

In this research, the systematic review was carried out in the first half of October 2021 using the following four databases: (1) ACL Anthology database, the main base for NLP, (2) IEEExplore, (3) ACM Digital Library, and (4) Scopus.

2.1 Protocol

The review aims to identify how syntactic information is utilized for the automatic detection of fake news and, ultimately, whether using syntactic information yields improvements. To better understand the context of use of syntactic information, it is also important to comprehend the datasets and classifiers involved in this process, as they significantly influence the overall results. With that in mind, our research questions were the following:

- What syntactic representations are used for fake news detection?
- What data extraction techniques are commonly used for fake news detection?
- What corpora or datasets are typically used for fake news detection?
- What classifiers are commonly applied in the detection of fake news?
- What are the effects of incorporating syntactic information into the results of fake news detection?
- What are the metrics used to measure the performance of fake news detection models?

During the search, only articles written in Portuguese or English were included, with no other filtering applied to the search engine. These were fetched through the search string *"fake news" and ("automatic detection" or "classification")* in their title or abstract, and *(syntactic or syntax)* in the corresponding text. The only exception was when searching the ACL Anthology database, in which *"fake news" and ("automatic detection" OR "classification")* was also applied to full texts, due to limitations in its search engine.

We decided to exclusively use the term "fake news", instead of similar terms such as hoaxes and misinformation, because it aligns more closely with the article's purpose. As defined by Zannettou *et al.* [2019], the category of fake news comprises the most severe type of false information found on the Web.

As a restriction, retrieved articles should have undergone a peer-review process, being also filtered according to some predefined inclusion and exclusion criteria. To be included in the reviewing process, articles should meet all inclusion criteria. The compliance with any of the exclusion criteria would rule the article out of further consideration. As a final step, selected articles were then fully read to answer the research questions.

To be considered in the survey, the article should then comply with the following inclusion criteria:

- They should satisfy the search string; and
- They should be fully available for download at the digital library.

The exclusion criteria for the retrieved articles were:

- Articles that do not report on the use of machine learning;
- Articles that do not present any validation method;
- Articles not related to the objective of the research
- Articles that were not peer-reviewed.

The "Articles not related to the objective" criterion excludes articles that, although containing the terms "fake news" and "syntax" or "syntactic", do not deal with the automatic detection of fake news with the aid of syntactic information. For example, while Cignarella *et al.* [2020] employ the usage of syntactic information, it is about irony detection and the term "fake news" is only present in the Acknowledgements section. The work from Kuzmin *et al.* [2020], on the other hand, is about fake news, but the term "syntactic" is only present in the related work. Since these articles are not inline with this research's objectives, they were excluded from our analysis.

3 Execution

The search in the four selected search engines returned a total of 230 articles, which were reduced to 20 articles after applying the inclusion and exclusion criteria. To analyze them, we initially reviewed their titles and abstracts. In cases where there was uncertainty, we proceeded to the analysis of their full text. This search and filtering process was done just once.

The vast majority of the articles were excluded by the "Articles not related to the objective of the research" criterion, since most of them only contained the word "syntax" or "syntactic" in their "related works" or "references" section, not being addressed by the research they described. Others contained the term "fake news", but that was not the focus of their studies.

Figure 1 shows the number of selected articles from each database, whereas Figure 2 presents the number of selected articles by year. Table 1 contains relevant information from the selected articles, such as the used corpus, and a column that summarizes the impact of the syntactic information on the performance metrics that are categorized by the following criteria. In this table, *Nothing* means it does not improve, or even degrades overall metrics. *Low* means the syntactic information marginally improved overall metrics. *Relevant* means the improvement is noticeable. *Not detailed* means the impact of syntactic information alone was not mentioned.

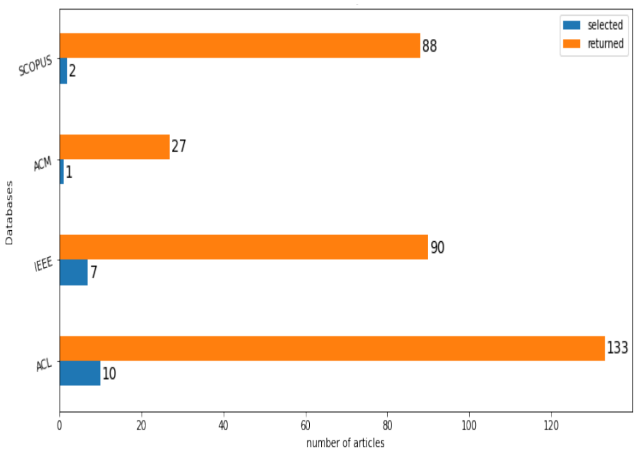

Figure 1. Number of selected articles from each database.

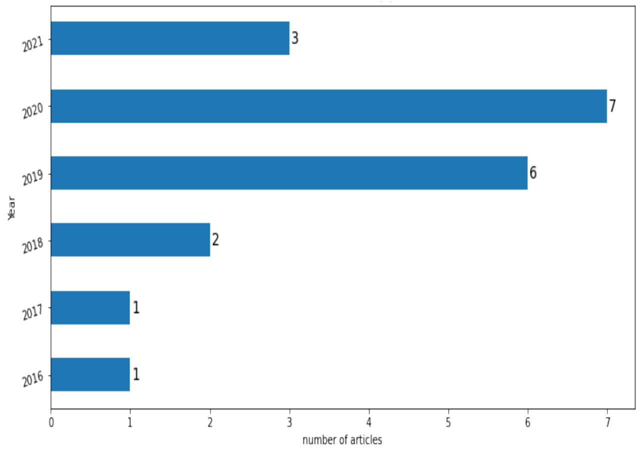

Figure 2. Number of selected article by publication year.

4 Results

This section aims to answer the research questions defined in the systematic review protocol by analyzing the 20 collected articles.

4.1 What syntactic representations are used for fake news detection?

Most articles use only morphosyntactic information at the word level, which is a very simple syntactic representation that does not adequately capture the syntactical relation between the words, this representation takes the form of parts of speech (PoS) [De Sarkar *et al.*, 2018; Vincze and Szabó, 2020; Rubin *et al.*, 2016; Hossain *et al.*, 2020; Joo and Hwang, 2019; Paixão *et al.*, 2020; Reis *et al.*, 2019; Agarwal and Dixit, 2020; Sabeeh *et al.*, 2019; Coste *et al.*, 2020; Choudhary and Arora, 2021; Zhou *et al.*, 2020]. Hossain *et al.* [2020] define this representation as *"shallow syntax"*, Vincze and Szabó [2020] perform the extraction of syntactic information only for some specific grammatical classes. Kapusta *et al.* [2021] consider Char count, Word count, Title word count, STOP word count, Upper case word as syntactic information.

Hossain *et al.* [2020] combine parts of speech with n-gram to better capture the relationships between words. Rubin *et al.* [2016] try Bidirectional Long short-term memory (BiLSTM) Based on PoS information to catch the syntactic information.

The remaining articles use more complex syntactical representations that explore the relationship between words. Some articles use context-free grammars (CFG) to model syntactic relations [Qiao *et al.*, 2020; Pérez-Rosas *et al.*, 2018; Han and Mehta, 2019; Zhou *et al.*, 2020]. Others explore the syntax through dependency trees [Volkova *et al.*, 2017; Nguyen *et al.*, 2019; Gupta *et al.*, 2021].

Figure 3 below shows the occurrence of each syntactic representation, considering the Kapusta *et al.* [2021] representation as part of the 'Others' category, and Vincze and Szabó [2020]; Hossain *et al.* [2020]; Rubin *et al.* [2016] as only PoS

Table 1. Summary of the selected articles

Article	Corpus	Syntactic information	Impact on the results
Gupta *et al*. [2021]	Proposed Corpus, evaluate on 6 Benchmarks	Novel deep neural framework Part-of-speech (PoS) Module Dependency Tree (DEP) Module	Relevant
Gupta *et al*. [2021]	Proposed Corpus, evaluate on 6 Benchmarks	Novel deep neural framework Part-of-speech (PoS) Module Dependency Tree (DEP) Module	Relevant
Choudhary and Arora [2021]	Random Political News BuzzFeed new	Char count; Word count; Title word count, STOP word count; Upper case word count;	Relevant
Kapusta *et al*. [2021]	Covid news	Part-Of-Speech (PoS)	Relevant
Qiao *et al*. [2020]	ISOT e LIAR	Part-of-speech (PoS); Probabilistic Context Free Grammar Parser	Relevant
Vincze and Szabó [2020]	Proposed Corpus	frequency of subjects and objects ; frequency of adverb; frequency of Coordination and Subordination	Low
Hossain *et al*. [2020]	Proposed Corpus	Part-of-speech (PoS)	Nothing
Paixão *et al*. [2020]	Fake.Br corpus	Part-of-speech (PoS)	Low
Agarwal and Dixit [2020]	Liar dataset	Part-of-speech (PoS)	Not detailed
Coste *et al*. [2020]	The Clickbait Challenge	Part-Of-Speech (PoS)	Relevant
Zhou *et al*. [2020]	PolitiFact and BuzzFeed	Part-Of-Speech (PoS) Probability Context Free Grammar (PCFG)	Relevant
Sabeeh *et al*. [2019]	Fever dataset	Part-of-speech (PoS)	Not detailed
Balwant [2019]	Liar dataset	Bidirectional LSTM Based on PoS	Relevant
Reis *et al*. [2019]	Buzz feed	Part-of-speech (PoS)	Not detailed
Han and Mehta [2019]	Rumor Twitter	Context free grammars (CFG) Encoded in TFIDF	Relevant
Nguyen *et al*. [2019]	(“hyper-partisan” vs. “not hyper-partisan”)	Ngram PoS tagging, Dependency sub-trees	Relevant
Joo and Hwang [2019]	SemEval 2019 task 4 dataset	Part-of-speech (PoS)	Relevant
Pérez-Rosas *et al*. [2018]	FakeNewsAMT, Celebrity	Context free grammars (CFG)	Relevant
De Sarkar *et al*. [2018]	Dataset com sátiras	Part-of-speech (PoS); SentiWordnet positive score; SentiWordnet negative score; SentiWordnet objective score; Named Entity IOB tag; Named Entity tag; Topmost Wordnet	Low
Volkova *et al*. [2017]	Twitter corpus	Part-of-speech (PoS) Dependency tree parses with SyntaxNe	Nothing
Rubin *et al*. [2016]	Proposed Corpus	Part-of-speech (PoS)	Relevant

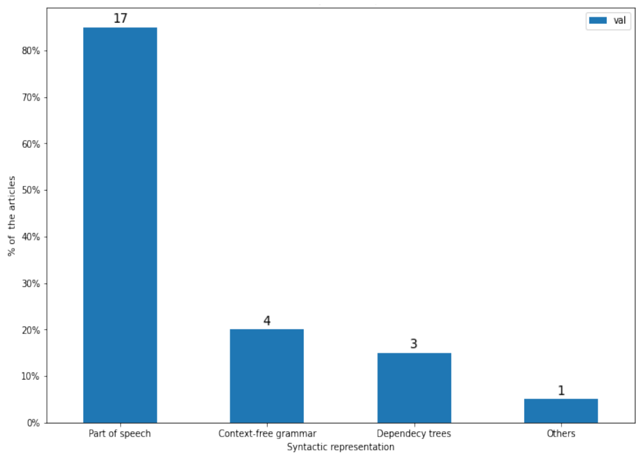

Figure 3. Occurrence of each syntactic representation.

4.2 What data extraction techniques are commonly used for fake news detection?

To extract parts of speech, a PoS Tagger is necessary, which varies across languages. Volkova *et al.* [2017] use a tool called SyntaxNe to extract dependency characteristics. For articles applying context-free grammar, a parser proposed by Klein and Manning [2003]. Articles that utilize dependency trees need a parser to extract the tree [Volkova *et al.*, 2017; Nguyen *et al.*, 2019; Gupta *et al.*, 2021].

There are several techniques for feature extraction, most studies adopt a straightforward method to extract features from PoS tags, such as the normalized frequency of each PoS tag [Paixão *et al.*, 2020; Reis *et al.*, 2019; Agarwal and Dixit, 2020; Sabeeh *et al.*, 2019; Qiao *et al.*, 2020; De Sarkar *et al.*, 2018; Rubin *et al.*, 2016; Hossain *et al.*, 2020; Joo and Hwang, 2019]. Some other studies employ more sophisticated methods. Balwant [2019] calculate the average word embedding for the same PoS tags, which is then passed to a bidirectional LSTM to obtain a vector representation of the news content. Kapusta *et al.* [2021] and Nguyen *et al.* [2019] create an n-gram model using PoS tags instead of words. Gupta *et al.* [2021] have a separate PoS Module which consists of an embedding layer followed by a BiLSTM and an attention layer.

Pérez-Rosas *et al.* [2018] extract a set of features derived from production rules based on CFG trees. Han and Mehta [2019] use TF-IDF to encode the syntactic information produced by context-free grammars. Nguyen *et al.* [2019] use text mining techniques to extract features from the syntactic tree of dependencies. Gupta *et al.* [2021] use a Transformer Encoder to capture features from the dependency tree.

4.3 What corpora or datasets are typically used for fake news detection?

Some articles propose their own datasets [Vincze and Szabó, 2020; Rubin *et al.*, 2016; Hossain *et al.*, 2020; Gupta *et al.*, 2021]. Volkova *et al.* [2017] performed their own annotations. There are a few datasets with a small amount of texts, such as [Vincze and Szabó, 2020; Rubin *et al.*, 2016] with 180 and 360 news articles respectively.

Different corpora usually present different definitions of fake news: deliberately false news [Qiao *et al.*, 2020; Vincze and Szabó, 2020], satires [Rubin *et al.*, 2016], clickbait news [Vincze and Szabó, 2020; Coste *et al.*, 2020]. hyper-partisans news [Nguyen *et al.*, 2019; Joo and Hwang, 2019], facts and fictions [Volkova *et al.*, 2017] and rumors [Han and Mehta, 2019]

Most corpora are in English, except for the ones introduced by Vincze and Szabó [2020]; Hossain *et al.* [2020]; Paixão *et al.* [2020] which are in Hungarian, Bengali and Portuguese respectively

4.4 What classifiers are commonly applied in the detection of fake news?

For classification, various classifiers and techniques are used, as illustrated in Figure 4.

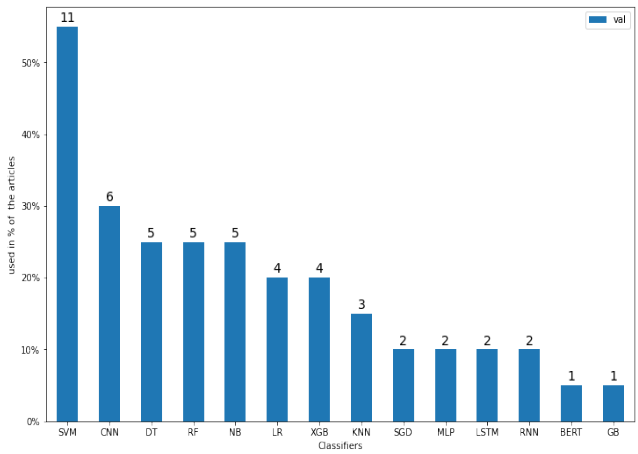

Figure 4. Summary of used classifiers across all selected articles.

SVM is perhaps so commonly used due to its popularity in the NLP field. SVM works well with unstructured and semi-structured data like texts. Also, SVM is flexible due to its different kernels. Those reasons might make SVM a good baseline classifier for more complex ones.

Qiao *et al.* [2020] tried several algorithms such as Stochastic Gradient Descent (SGD), Support Vector Machines, Linear Support Vector Machines (LSVM), K-Nearest Neighbour(KNN), and Decision Trees (DT). LSVM turned out to be the best. De Sarkar *et al.* [2018] and Gupta *et al.* [2021], after generating complex embeddings, use a multilayer perceptron (MLP) for the final classification. . Pérez-Rosas *et al.* [2018], Rubin *et al.* [2016] and Nguyen *et al.* [2019] combine all the features into a linear SVM classifier. Vincze and Szabó [2020] and Coste *et al.* [2020] use random forests.

Hossain *et al.* [2020] try classic approaches such as SVM, Random Forest (RF), and Logistic Regression (LR), along of with more recent ones, such as Convolutional Neural Networks (CNN), LSTM, and Bidirectional Encoder Representations from Transformers (BERT). Surprisingly, SVM outperforms all these recent approaches. Conversely, Volkova *et al.* [2017] and Vincze and Szabó [2020] obtained the best performance with CNN.

Joo and Hwang [2019] adopt the average weighted value of the style-based approach using XGBoost (XGB) and the content-based approach using CNN.. Balwant [2019] generate an embedding vector out of PoS tags that is the input to a softmax classifier. Paixão *et al*. [2020] test several classifiers, including Linear SVM, LR, Naive Bayes (NB), DT, RF, and XGB and CNN. In their research, CNN outperformed others by far. Reis *et al*. [2019] tested KNN, NB, RF, SVM with RBF kernel and XGB. The best results were obtained by RF and XGB classifiers.

Agarwal and Dixit [2020] test multiple classifiers, and LSTM shows higher performance than SVM, CNN, KNN, and NB. Sabeeh *et al*. [2019] use DT and SVM, with a better performance for the latter. Han and Mehta [2019] explore CNN and Recurrent Neural Network (RNN) hybrid model, among others such as Decision Tree and Gradient Boosting (GB), SGD, SVM, NB, the hybrid model presents the best performance. Zhou *et al*. [2020] test the following classifiers LR, NB, SVM (with linear kernel), RF, and XGB. The best results were obtained with XGB, SVM and RF. Kapusta *et al*. [2021] classify fake news with a Decision Tree.

4.5 What are the effects of incorporating syntactic information into the results of fake news detection?

Pérez-Rosas *et al*. [2018] conclude that it is worthwhile to look at the lexical, syntactic, and semantic levels. Rubin *et al*. [2016] discovered that individual textual features of PoS and punctuation marks are highly indicative of the presence of satire, producing a detection improvement of 5%. The sophisticated architecture with syntactic features proposed by Gupta *et al*. [2021] presented an increase of up to 3 % in the performance. Kapusta *et al*. [2021] noticed an accuracy increase by 3 to 4 %, where the average accuracies for syntactic, readability, and combined features are 84.12%, 77.67%, and 84.52%, respectively.

Nguyen *et al*. [2019] identified that n-grams in combination with dependency sub-trees as features have a positive impact on the performance of the classifier. Balwant [2019] indicated that the proposed bidirectional LSTM exploits PoS tags information of news content and outperforms state-of-art architectures that do not exploit syntactic information. Reis *et al*. [2019] obtained good results using PoS tags features.

Han and Mehta [2019] showed that the set of features based on Context-Free Grammar had potential predictive power. Zhou *et al*. [2020] saw no improvement with shallow syntax, but deep syntax-level features (CFGs) and features at lexicon-level (BOWs) outperform the others. Kapusta *et al*. [2021] concluded that morphological analysis can be applied to fake news classification.

Studies try several classifiers and, in general, SVM presented good results for most studies. Some studies, after applying and combining with multiple techniques, use a neural network classifier with a hidden layer or just softmax at end of the classification process.

Some studies did not see any improvement using syntactic features [Hossain *et al*., 2020; Volkova *et al*., 2017]. Others could only see an improvement in performance with syntactic features when they are combined with others [Vincze and Szabó, 2020; Joo and Hwang, 2019; Paixão *et al*., 2020; Qiao *et al*., 2020]. Some studies included syntactic features in their final set but did not mention their separate impact [Paixão *et al*., 2020; Reis *et al*., 2019; Agarwal and Dixit, 2020].

Figure 5 summarizes the impact of the syntactic information on the performance metrics. In this Figure 5, *Nothing* means no improvement, or even a degradation in overall metrics. *Low* means the syntactic information marginally improved overall metrics. *Relevant* means the improvement is noticeable. *Not detailed* means the impact of syntactic information alone was not mentioned.

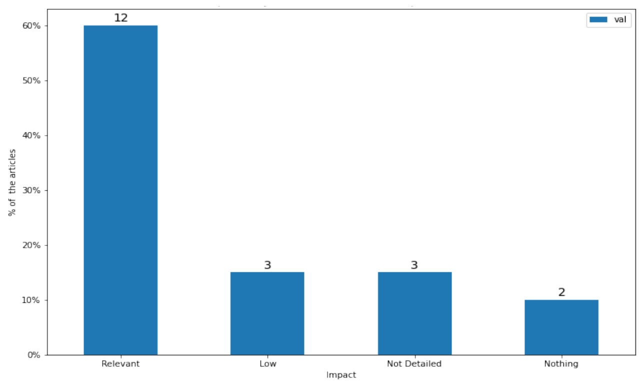

Figure 5. Impact of syntactic information on the performance.

4.6 What are the metrics used to measure the performance of fake news detection models?

Accuracy, precision, recall, and F1 score are the most common metrics. Few studies use AUC (area under the ROC curve) [Volkova *et al*., 2017; Reis *et al*., 2019; Han and Mehta, 2019]. Accuracy and F1 are generally the main metrics for comparison across studies.

When dealing with corpora containing more than two classes, researchers often compute the macro-averaged F1 score (*e.g.* Gupta *et al*. [2021]; Reis *et al*. [2019]), which is the arithmetic mean of F1 scores across all classes. Figure 6 shows the results of an analysis of performance metrics used in selected articles. The y-axis of the figure represents the percentage of articles that used a particular metric. The number on top of each bar represents the absolute count, which is the number of articles that used a particular metric.

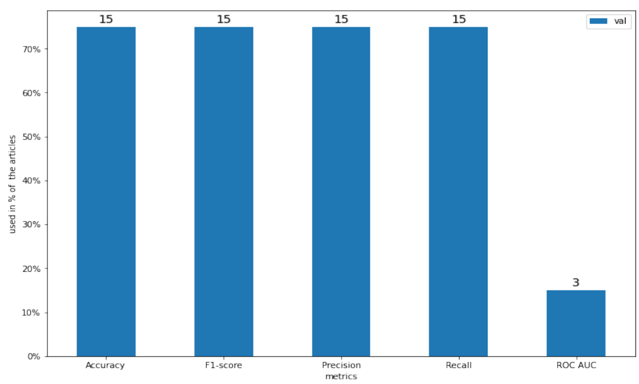

Figure 6. Relative and Absolute count of articles that used each metric.

Table 2. Summary of language models and representations

Language Model and Representation	Articles
bag of words relative or/and absolute frequency	[Qiao *et al*., 2020; Vincze and Szabó, 2020; Paixão *et al*., 2020] [Reis *et al*., 2019; Agarwal and Dixit, 2020; Zhou *et al*., 2020]
n-gram of words	[Qiao *et al*., 2020; Paixão *et al*., 2020; Hossain *et al*., 2020] [Nguyen *et al*., 2019; Reis *et al*., 2019; Agarwal and Dixit, 2020]
n-gram of characters	[Qiao *et al*., 2020; Hossain *et al*., 2020]
TF-IDF	[Rubin *et al*., 2016; Paixão *et al*., 2020; Agarwal and Dixit, 2020] [Kapusta *et al*., 2021]
n-gram encoded as TF-IDF	[Pérez-Rosas *et al*., 2018; Han and Mehta, 2019; Joo and Hwang, 2019]
word embedding with word2vec	[Joo and Hwang, 2019; Agarwal and Dixit, 2020; Sabeeh *et al*., 2019] [Balwant, 2019]
word embedding with glove	[De Sarkar *et al*., 2018; Volkova *et al*., 2017]
word embedding with TextToSequence Keras API	[Reis *et al*., 2019]
BERT	[Gupta *et al*., 2021]

4.7 Language models

Although describing the language models adopted in each research was not initially defined as a research question in our protocol, we found this information relevant after analyzing the selected articles, so we decided to add it here, as an extra feature. in this regard, different language models are used alongside syntactic information, with n-gram and bag of words being the most popular ones. TF-IDF encoding is highly used to generate features. Also, Word embedding representations are popular, mostly obtained via Glove and Word2Vec techniques.

There are some combinations of these techniques, such as bigrams derived from the bag of words encoded as TF-IDF in [Pérez-Rosas *et al*., 2018]. Coste *et al*. [2020] rely only on features that are language independent and therefore did not explore language models. Table 2 describes the language models and their representations more granularly.

5 Discussion

Most studies use simple syntactic information at word-level. They are usually all part of speech tags or specific PoS information such as the frequency of adverbs. To extract features from this shallow syntax, studies use different techniques, from a simple normalized frequency of each PoS tag to a complex architecture with BiLSTM to encode PoS information.

Studies that used this kind of information had very marginal improvements when combined with other techniques or no improvement, except for Rubin *et al*. [2016] that achieved a significant improvement.

Other studies choose more complex syntactic structures, to better capture the syntax information at the sentence level. The most used ones are Dependency Tree and Context-Free grammars. There are also different techniques to extract features from them, such as Transformer Encoder, Text mining and TFIDF.

In general, studies that exploit deep syntax present a more robust performance, with highlights to Gupta *et al*. [2021] with a solution using Dependency Tree and PoS beats several benchmarks in the literature. Volkova *et al*. [2017] is the only exception to this rule, contrary to their expectations, syntax and grammar features worsen performance, they hypothesized that Twitter could be the cause for this result, given the fact its messages are usually shorter, noisier, and difficult to parse.

Although just a few, studies with Corpora in languages other than English present very promising results. Vincze and Szabó [2020], Hossain *et al*. [2020] and Paixão *et al*. [2020] with Corpora in Hungarian, Bengali, and Portuguese, respectively, obtained, at the best, marginal increases in the performance. This may be due to the parses used in each language.

6 Conclusion

This article presented a systematic review of the automatic detection of fake news using syntactic information. Through this study, we aimed to better understand the usage of syntactic information in the context of fake news. We formulated some research questions regarding the most used techniques and their overall impact on performance. These research questions were answered based on 20 collected articles in the literature.

In general, shallow syntax representations, such as Parts of Speech, only marginally increase the performance. However, deep syntactic representations, such as context-free grammars or syntactic trees of dependencies, presented more promising results. They were able to improve overall performance when combined with other techniques or even outperform state-of-the-art models in some cases.

As a limitation of this research, our results heavily depend on the scope and accuracy of the search engines used to retrieve the articles. We attempted to mitigate this problem by using broader search terms to increase recall, albeit at the cost of reduced precision. This approach, however, is not guaranteed to have led us to a high recall rate. In addition to that, we want to note that our analysis is limited to studies published until the first half of October 2021. Any research or findings after this period are not included in our review.

Still, we believe this study to be particularly relevant given the thoroughness of its evaluation regarding the utility of syntactic information in the automatic detection of fake news. Our findings suggest that using syntactic information, especially in combination with other features, is a promising approach for improving performance in fake news detection. We recommend future studies to explore the techniques presented in this study and investigate other syntactic representations, such as constituency-based parse trees for example. Additionally, we suggest exploring the use of semantic features, which we did not cover in this study, as a potential way to improve performance further.

Declarations

Acknowledgements

The authors of this work would like to thank the Center for Artificial Intelligence (C4AI-USP) and the support from the São Paulo Research Foundation (FAPESP grant #2019/07665-4) and from the IBM Corporation.

Authors' Contributions

Competing interests

The authors declare that they have no competing interests.

Availability of data and materials

Data can be made available upon request.

References

Agarwal, A. and Dixit, A. (2020). Fake news detection: An ensemble learning approach. In *2020 4th International Conference on Intelligent Computing and Control Systems (ICICCS)*, pages 1178–1183, Madurai, India. IEEE. DOI: 10.1109/ICICCS48265.2020.9121030.

Allcott, H. and Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of Economic Perspectives*, 31(2):211–36. DOI: 10.1257/jep.31.2.211.

Balwant, M. K. (2019). Bidirectional lstm based on pos tags and cnn architecture for fake news detection. In *2019 10th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*, pages 1–6, Kanpur, India. IEEE. DOI: 10.1109/ICCCNT45670.2019.8944460.

Choudhary, A. and Arora, A. (2021). Linguistic feature based learning model for fake news detection and classification. *Expert Systems with Applications*, 169:114–171. DOI: 10.1016/j.eswa.2020.114171.

Cignarella, A. T., Basile, V., Sanguinetti, M., Bosco, C., Rosso, P., and Benamara, F. (2020). Multilingual irony detection with dependency syntax and neural models. In Scott, D., Bel, N., and Zong, C., editors, *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1346–1358, Barcelona, Spain (Online). International Committee on Computational Linguistics. DOI: 10.18653/v1/2020.coling-main.116.

Coste, C. I., Bufnea, D., and Niculescu, V. (2020). A new language independent strategy for clickbait detection. In *2020 International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, pages 1–6, Split, Croatia. IEEE. DOI: 10.23919/SoftCOM50211.2020.9238342.

De Sarkar, S., Yang, F., and Mukherjee, A. (2018). Attending sentences to detect satirical fake news. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3371–3380, Santa Fe, New Mexico, USA. Association for Computational Linguistics. Available at: `https://aclanthology.org/C18-1285/`.

Gupta, S., Singh, P., Sundriyal, M., Akhtar, M. S., and Chakraborty, T. (2021). Lesa: Linguistic encapsulation and semantic amalgamation based generalised claim detection from online content. DOI: 10.48550/arXiv.2101.11891.

Han, W. and Mehta, V. (2019). Fake news detection in social networks using machine learning and deep learning: Performance evaluation. In *2019 IEEE International Conference on Industrial Internet (ICII)*, pages 375–380, Orlando, USA. IEEE. DOI: 10.1109/ICII.2019.00070.

Hossain, M. Z., Rahman, M. A., Islam, M. S., and Kar, S. (2020). BanFakeNews: A dataset for detecting fake news in Bangla. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2862–2871, Marseille, France. European Language Resources Association. Available at: `https://aclanthology.org/2020.lrec-1.349`.

Joo, Y. and Hwang, I. (2019). Steve Martin at SemEval-2019 task 4: Ensemble learning model for detecting hyperpartisan news. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 990–994, Minneapolis, Minnesota, USA. Association for Computational Linguistics. DOI: 10.18653/v1/S19-2171.

Kapusta, J., Drlik, M., and Munk, M. (2021). Using of n-grams from morphological tags for fake news classification. *PeerJ Computer Science*, 7:1–27. DOI: 10.7717/PEERJ-CS.624.

Kitchenham, B. (2004). Procedures for performing systematic reviews. *Keele, UK, Keele University*, 33(2004):1–26. Available at:`https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=29890a936639862f45cb9a987dd599dce9759bf5`.

Klein, D. and Manning, C. D. (2003). Accurate unlexicalized parsing. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, pages 423–430, Sapporo, Japan. Association for Computational Linguistics. DOI: 10.3115/1075096.1075150.

Kuzmin, G., Larionov, D., Pisarevskaya, D., and Smirnov, I. (2020). Fake news detection for the Russian language. In Aker, A. and Zubiaga, A., editors, *Proceedings of the 3rd International Workshop on Rumours and Deception in Social Media (RDSM)*, pages 45–57, Barcelona, Spain (Online). Association for Computational Linguistics. Available at:`https://aclanthology.org/2020.rdsm-1.5`.

Lazer, D. M. J., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., Schudson, M., Sloman, S. A., Sunstein, C. R., Thorson, E. A., Watts, D. J., and Zittrain, J. L. (2018). The science of fake news. *Science*, 359(6380):1094–1096. DOI: 10.1126/science.aao2998.

Moscadelli, A., Albora, G., Biamonte, M. A., Giorgetti, D., Innocenzio, M., Paoli, S., Lorini, C., Bonanni, P., and Bonaccorsi, G. (2020). Fake news and covid-19 in italy: Results of a quantitative observational study. *International Journal of Environmental Research and Public Health*, 17. DOI: 10.3390/ijerph17165850.

Moschitti, A. (2006). Efficient convolution kernels for dependency and constituent syntactic trees. In Fürnkranz, J., Scheffer, T., and Spiliopoulou, M., editors, *Machine Learning: ECML 2006*, pages 318–329, Berlin, Heidelberg. Springer Berlin Heidelberg. DOI: $10.1007/11871842_32$.

Nguyen, D.-V., Dang, T., and Nguyen, N. (2019). NLP@UIT at SemEval-2019 task 4: The paparazzo hyperpartisan news detector. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 971–975, Minneapolis, Minnesota, USA. Association for Computational Linguistics. DOI: 10.18653/v1/S19-2167.

Oshikawa, R., Qian, J., and Wang, W. Y. (2020). A survey on natural language processing for fake news detection. DOI: 10.48550/arXiv.1811.00770.

Paixão, M., Lima, R., and Espinasse, B. (2020). Fake news classification and topic modeling in brazilian portuguese. In *2020 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*, pages 427–432, Melbourne, Australia. IEEE. DOI: 10.1109/WIIAT50758.2020.00063.

Pérez-Rosas, V., Kleinberg, B., Lefevre, A., and Mihalcea, R. (2018). Automatic detection of fake news. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3391–3401, Santa Fe, New Mexico, USA. Association for Computational Linguistics. Available at:`https://aclanthology.org/C18-1287/`.

Pisarevskaya, D. (2017). Deception detection in news reports in the Russian language: Lexics and discourse. In *Proceedings of the 2017 EMNLP Workshop: Natural Language Processing meets Journalism*, pages 74–79, Copenhagen, Denmark. Association for Computational Linguistics. DOI: 10.18653/v1/W17-4213.

Qiao, Y., Wiechmann, D., and Kerz, E. (2020). A language-based approach to fake news detection through interpretable features and BRNN. In *Proceedings of the 3rd International Workshop on Rumours and Deception in Social Media (RDSM)*, pages 14–31, Barcelona, Spain (Online). Association for Computational Linguistics. Avalable at: `https://aclanthology.org/2020.rdsm-1.2/`.

Reis, J. C. S., Correia, A., Murai, F., Veloso, A., and Benevenuto, F. (2019). Supervised learning for fake news detection. *IEEE Intelligent Systems*, 34(2):76–81. DOI: 10.1109/MIS.2019.2899143.

Rubin, V., Conroy, N., Chen, Y., and Cornwell, S. (2016). Fake news or truth? using satirical cues to detect potentially misleading news. In *Proceedings of the Second Workshop on Computational Approaches to Deception Detection*, pages 7–17, San Diego, California. Association for Computational Linguistics. DOI: 10.18653/v1/W16-0802.

Rubin, V. L., Chen, Y., and Conroy, N. J. (2015). Deception detection for news: Three types of fakes. In *Proceedings of the 78th ASIS&T Annual Meeting: Information Science with Impact: Research in and for the Community*, ASIST '15, USA. American Society for Information Science. DOI: 10.1002/pra2.2015.145052010083.

Sabeeh, V., Zohdy, M., and Bashaireh, R. A. (2019). Enhancing the fake news detection by applying effective feature selection based on semantic sources. In *2019 International Conference on Computational Science and Computational Intelligence (CSCI)*, pages 1365–1370, Las Vegas, USA. IEEE. DOI: 10.1109/CSCI49370.2019.00255.

Santos, R., Pedro, G., Leal, S., Vale, O., Pardo, T., Bontcheva, K., and Scarton, C. (2020). Measuring the impact of readability features in fake news detection. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1404–1413, Marseille, France. European Language Resources Association. Available at: `https://aclanthology.org/2020.lrec-1.176/`.

Shu, K., Sliva, A., Wang, S., Tang, J., and Liu, H. (2017a). Fake news detection on social media: A data mining perspective. *SIGKDD Explor. Newsl.*, 19(1):22–36. DOI: 10.1145/3137597.3137600.

Shu, K., Sliva, A., Wang, S., Tang, J., and Liu, H. (2017b). Fake news detection on social media: A data mining perspective. *SIGKDD Explor. Newsl.*, 19(1):22–36. DOI: 10.1145/3137597.3137600.

Silva, R. M., Santos, R. L., Almeida, T. A., and Pardo, T. A. (2020). Towards automatically filtering fake news in portuguese. *Expert Systems with Applications*, 146:113199. DOI: 10.1016/j.eswa.2020.113199.

Vincze, V. and Szabó, M. K. (2020). Automatic detection of Hungarian clickbait and entertaining fake news. In *Proceedings of the 3rd International Workshop on Rumours and Deception in Social Media (RDSM)*, pages 58–69, Barcelona, Spain (Online). Association for Computational Linguistics. Available at:`https://aclanthology.org/2020.rdsm-1.6/`.

Volkova, S., Shaffer, K., Jang, J. Y., and Hodas, N. (2017). Separating facts from fiction: Linguistic models to classify suspicious and trusted news posts on Twitter. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 647–653, Vancouver, Canada. Association for Computational Linguistics. DOI: 10.18653/v1/P17-2102.

Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380):1146–1151. DOI: 10.1126/science.aap9559.

Zannettou, S., Sirivianos, M., Blackburn, J., and Kourtellis, N. (2019). The web of false information: Rumors, fake news, hoaxes, clickbait, and various other shenanigans. *J. Data and Information Quality*, 11(3). DOI: 10.1145/3309699.

Zewdu, A. and Yitagesu, B. (2022). Part of speech tagging: a systematic review of deep learning and machine learning approaches. *Journal of Big Data*, 9. DOI: 10.1186/s40537-022-00561-y.

Zhou, X., Jain, A., Phoha, V. V., and Zafarani, R. (2020). Fake news early detection: A theory-driven model. *Digital Threats: Research and Practice*, 1(2). DOI: 10.1145/3377478.

Zhou, X. and Zafarani, R. (2020). A survey of fake news: Fundamental theories, detection methods, and opportunities. *ACM Comput. Surv.*, 53(5). DOI: 10.1145/3395046.

Zhou, X., Zafarani, R., Shu, K., and Liu, H. (2019). Fake news: Fundamental theories, detection strategies and challenges. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, WSDM '19, page 836–837, New York, NY, USA. Association for Computing Machinery. DOI: 10.1145/3289600.3291382.