



USING NATURAL LANGUAGE PROCESSING IN DETECTING FAKE NEWS

Lappeenranta–Lahti University of Technology LUT

Bachelor's Programme in Software Engineering, Bachelor's thesis

2022

Vivian Künnap

Examiner: Professor Annika Wolff

ABSTRACT

Lappeenranta–Lahti University of Technology LUT

LUT School of Engineering Science

Software Engineering

Vivian Künnap

Using natural language processing in detecting fake news

Bachelor's thesis

2022

34 pages, 5 figures and 1 table

Examiner: Professor Annika Wolff

Keywords: natural language processing, fake news, deep learning, text classification

Fake news is not a novel concept but the scale of its spread and the damage it has and continues to cause is alarming. From the US presidential elections in 2016 to the COVID-19 pandemic and today, fake news has been circulating in news media corrupting the public opinion. Fake news alters democratic discussions polarizing people's opinions sowing distrust in national institutions and setting different groups against each other. It is a threat to democracy and national security. It is crucial to prevent fake news from spreading and one solution is to create an automatic fake news detection system. A solution is researched using natural language processing (NLP) tasks, namely text classification. NLP is a type of artificial intelligence that is essentially taught to understand human language. Using thematic analysis, the main steps and techniques of fake news detection models are described and through a comparative analysis the state-of-the-art models are distinguished. And while there are many potential fake news detection models for English there is not much variety for other languages. So, it is additionally analysed if these benchmark models can be implemented for Finnish language as well.

TIIVISTELMÄ

Lappeenrannan–Lahden teknillinen yliopisto LUT

LUT Teknis-luonnontieteellinen

Tietotekniikka

Vivian Künnap

Valeutisten havaitseminen käyttäen luonnollisen kielen käsittelyjärjestelmiä

Tietotekniikan kandidaatintyö

2022

34 sivua, 5 kuvaa ja 1 taulukko

Tarkastaja: Professori Annika Wolff

Avainsanat: luonnollisen kielen käsittelyjärjestelmä, valeutiset, syväoppiminen, tekstin luokittelu

Valeutiset eivät ole uusi käsite, mutta niiden leviämisen laajuus ja niiden aiheuttamat vahingot ovat huolestuttavia. Yhdysvaltain presidentinvaaleista vuonna 2016 COVID-19-pandemiaan ja nykypäivään asti, uutismediassa on kiertänyt valeutisia, jotka muokkaavat yleisön mielipidettä. Valeutiset muuttavat demokraattista keskustelua polarisoimalla ihmisten mielipiteitä kylväen epäluottamusta kansallisiin instituutioihin ja asettaen erilaisia ryhmiä toisiaan vastaan. Se on uhka demokratialle ja kansalliselle turvallisuudelle. On tärkeää estää valeutisten leviäminen, ja yksi ratkaisu on luoda automaattinen valeutisten havaitsemisjärjestelmä. Ratkaisua tutkitaan käyttämällä luonnollisen kielen käsittelyn (NLP) tehtäviä, etenkin tekstin luokittelua. NLP on tekoälyn tyyppi, missä tietokone opetetaan ymmärtämään ihmisten kieltä. Temaattisen analyysin avulla kuvataan valeutisten havaitsemismallien päävaiheet sekä tekniikat, ja vertailevan analyysin avulla valikoidaan uusimmat ja onnistuneimmat mallit. Ja vaikka englannin kielellä on monia mahdollisia valeutisten havaitsemismalleja, muille kielille ei ole paljon valikoimaa. Lisäksi analysoidaan, voidaanko nämä mallit toteuttaa myös suomen kielelle.

SYMBOLS AND ABBREVIATIONS

Abbreviations

NLP	Natural Language Processing
TF	Term Frequency
TF-IDF	Term Frequency-Inverted Document Frequency
GloVe	Global Vectors for Word Representation
DNN	Deep Neural Network
CNN	Convolutional Neural Network
RNN	Recurrent Neural Network
LSTM	Long Short-Term Memory
CSI	Capture, Score, Integrate
BERT	Bidirectional Encoder Representations from Transformers

Table of contents

Abstract

Symbols and abbreviations

1	Introduction	7
1.1	Objective and boundary	8
1.2	Research questions and methods.....	8
1.3	Thesis structure	8
2	Background and related works	9
2.1	Fake news.....	9
2.1.1	Characteristics of fake news	9
2.1.2	News sites	10
2.1.3	Social media.....	11
2.2	Natural language processing	12
2.3	Related works.....	14
3	Model techniques.....	15
3.1	Dataset and pre-processing	15
3.2	Word embedding and feature extraction	16
3.3	Classification methods	18
3.3.1	CNN-based.....	18
3.3.2	LSTM-based	20
3.3.3	CSI.....	21
3.3.4	BERT-based.....	21
4	Comparison of results.....	23
4.1	Finnish fake news detection	25
5	Discussion and conclusions	26
5.1	Central conclusions to research questions.....	27
5.2	Research bias and reliability	28
5.3	Future outlook and prospects	28
5.4	Future works.....	29
	References.....	30

Figures

Figure 1: Illustration of a neuron as a linear function.

Figure 2: Visualization of data processing in detection methods.

Figure 3: Example of n-grams.

Figure 4: Visualization of how CNN works.

Figure 5: An LSTM cell.

Tables

Table 1: Highest accuracy scores.

1 Introduction

Fake news started to flood the real news during the 2016 presidential elections in the United States of America. This phenomenon was not completely novel, but the most concerning aspect was that fake news were presented as actual facts. Many people were quick to believe in most of these so-called facts presented by their presidential candidate rather than researching the truth themselves. The surge of fake news gathered new momentum again during the COVID-19 pandemic. A time when everyone was forced to isolate and normal life was put on hold, many became frustrated and angry with the government. These types of environments are perfect for fake news to thrive. Sowing distrust and resentment in a part of the public becomes a volatile system. This materialized on the 6th of January in 2021 in the US as the attack on the Capitol. Another life-threatening example is anti-vaxxers. In fake news, vaccines are demonised, and their side effects are either greatly exaggerated or falsified. This feeds on the readers fear and has led many parents to not vaccinating their children and therefore predisposing them to diseases that should not be appearing anymore. Now with the Russian attack in Ukraine a new flood is rising again. Fake news causes misinformation that has politically polarized many and is endangering the health of them and those around them. This continues to be the problem and is why this topic is paramount.

Natural language processing (NLP) is one possible solution to tackling this problem. NLP is a part of artificial intelligence where a computer is taught how to understand text and spoken language. A suitable program using NLP could read different news sources on the internet and review how reliable they are. Reliability could be assessed for instance by going over the references in the text, doing a background check on the author and finding the intention behind the text. While there are already manual fake news detection websites like PolitiFact, an automated detector would scale larger, and the results could be available instantly. Creating a functional fake news detecting program would help people sort through the false and the correct. This would let them make choices and base their opinions on facts.

1.1 Objective and boundary

This bachelor's thesis is conducted for Lappeenranta-Lahti University of Technology LUT. The objective of this research is to find the state-of-the-art models for fake news detection using NLP techniques. Additional purpose is to evaluate if these solutions can be used for Finnish language platforms, i.e., news sites and news spread through social media. Thus, I set the boundary for research papers and other thesis to be from the last five years to find the latest and most advanced solutions. For source criticism, research papers are looked up from Elsevier's ScienceDirect.com, ACM digital libraries, IEEE Xplore and lastly Google Scholar.

1.2 Research questions and methods

The following research questions are set:

- How is NLP used to detect fake news online?
- What types of NLP have been used for this purpose and what have been the results?
- Can these be used for Finnish language platforms?

I will attempt to answer these questions through a thematic analysis and finding frequently occurring techniques. For the second question, some of the best models are gathered and compared. I will answer the third question by analysing found solutions and additionally researching if there already exists fake news detection NLPs for the Finnish language.

1.3 Thesis structure

Section 2 contains background and related works. In section 3, the main themes of fake news detectors are described. In section 4 the models with the most accurate accuracy scores are compared. Section 5 concludes the thesis and presents some discussion and proposes points for future works.

2 Background and related works

In this section, definitions for the essential terms used throughout this thesis are presented namely fake news and NLP. For fake news, commonly observed characteristics as well as some typical instances of fake news in news websites and social media are described.

2.1 Fake news

Zhou and Zafarani (2020) define fake news as intentionally false news published by a news outlet. Fake news' objective is to mislead and misinform the public in a malicious intent. This definition requires the authenticity of the news to be false. Meaning that using the term for articles that do not agree with an individual's or group's opinions or viewpoints is wrong. A news item can be partly or entirely false. Additionally, true news can be false when their subsequent events cause it to become obsolete. (Zhou & Zafarani 2020). Traditional and established news sites and media adheres to journalistic standards and ethics which their journalists are held accountable for, fake news on the other hand does not (Watson 2018).

Although humans are quite bad at recognizing fake news (Edkins 2016), in some cases, people unconsciously want to believe fake news. Shu et al. (2017) has presented two psychological factors to why people believe fake news: naivety and confirmation bias. People tend to believe that their perceptions are the correct ones and everyone else is considered irrational and ignorant. People also prefer to receive information that confirms their bias. This gives them a positive emotional feedback where they get to feel that they are in the right and everyone else is in the wrong.

2.1.1 Characteristics of fake news

To detect fake news, it is important to list some of their common characteristics. Fake news usually has sensationalist headlines, also called click baiting. Watson (2018) also claims that fake news is written by opportunists who are seeking revenue from their titles or want to manipulate political beliefs. As the intent of fake news is typically to alter public opinion they need to be spread vastly. Click bait titles are a common method to gain a lot of attention

producing revenue for the journalists and growing attention in the community making people talk more about the subject.

Fake news also uses provocative language. If the topic of the article fits their agenda, their statements tend to be exaggerated. If on the other hand, the topic is embarrassing or against their agenda, the statements tend to be either dismissive or not written about to begin with. Fake news articles also include more capitalisation, exclamation and question marks and less diversity in the vocabulary (Yang et al. 2022). Capitalisation and exclamation marks serve to bring more attention, question marks are used in rhetorical questions to guide readers' opinions and simple vocabulary guarantees that all readers understand the text. According to Yang et al. (2022), the images in fake news tend to be irrelevant or of low resolution.

Ruchansky, Seo and Liu (2017) have proposed a model that signals out three components of fake news: text, response, and source. Text is compared with the headline to see if their language matches. The type of response the article is meant to illicit. Usually, fake news articles provoke strong feelings in readers, e.g., anger or frustration. This is produced by using inflammatory and goading language. Finally, the source refers to the profile of the journalist and the credibility of the news media as well as the credibility of used sources and references used in the text.

In conclusion, fake news can be characterized by scandalous titles, provocative language that incites a strong emotional reaction, exaggerative or on the other hand disparaging statements, multitude of capitalisation and punctuation marks, images out of context or of low quality, the statements are not backed up by credible sources, and the journalist or the news site is not trustworthy.

2.1.2 News sites

Fake news articles published in actual news sites can be especially harmful since those are usually considered trustworthy sources. Generally, there are specific websites that regularly publish fake news. These sites have a distinct readership who harbour a certain set of values which these news cater to. Some of these sites use methods similar to phishing attacks for instance, by mimicking the URL of the true news sites. ABCnews.com.co is an example of

a phishing website, where the URL, logo, and design mock those of the real ABC.com news site.

There are some websites that determine if a news site is credible or not for instance MediaRank. MediaRank's algorithm analyses the website on four properties: peer reputation, reporting bias/breadth, bottom line financial pressure, and popularity (Ye & Skiena 2019). Based on these the readers can decide for themselves if the site is reliable or not. Although, in this case the readers are probably already careful about the veracity of the news they consume and less likely to believe fake news to begin with.

2.1.3 Social media

Social media is open for everyone's opinions and viewpoints. Anyone can post almost anything and everyone else is free to share these posts. Social media is also used increasingly more as a source of information, which means that its importance for journalists and news medias is also increasing. This leads to a substantial amount of clickbait articles rotating around and gaining considerable attention and hubbub. Social media posts are also abundantly difficult to moderate since the amount of information posted there is infinite and always growing.

Fake news in social media can be observed as either news sites posting links to their articles or creating news content specifically tailored to social media platforms. For example, YLE creates short vlog-style videos of news items to Instagram. In these cases, the aforementioned characteristics apply here. Fake news can also be presented as posts and comments written by civilians on their own accounts.

Some prevailing characteristics of fake news spread in social media include malicious accounts either trolls or bots, whose purpose is to share propaganda through clickbait articles and polarizing fake news (Shu et al. 2017). Bots are social media accounts controlled by computer algorithms to interact with other accounts and produce content (Ferrara et al. 2016). Bots become malicious when they produce and share fake news to distort public opinion and alter their behaviour. According to a study done by Bessi and Ferrara (2016), democratic discussion on social media during the 2016 presidential elections in the US was greatly affected and altered by the content created by bots. Trolls on the other hand are real

humans who aim to trigger negative emotional response in the community and undermine the credibility of real news sources with ridicule.

Another characteristic provided by Shu et al. (2017) is the echo chamber effect. Social media feed's algorithms show users similar type of content they like. Hence, users find more and more posts about the topics they are interested in and share the same opinions and viewpoints with. Users are exposed to their favoured narratives and form like-minded groups further polarizing their opinions. This leads to the echo chamber effect where people believe fake news catered to them due to peers believing them and the frequency of which they have seen these posts.

2.2 Natural language processing

Natural language processing (NLP) is a subcategory of artificial intelligence where a computer is thought to analyse, process, and produce natural language, i.e., human spoken language (IBM 2020). The aim of NLP is to produce an algorithm that understands natural language as well as the contextual background knowledge behind the data. A few everyday examples of NLP are speech-to-text dictation software, digital assistants, chatbots and GPS speech directions.

Creating an NLP is challenging due to our language being so complex and full of ambiguities. Words can have multiple different meanings and same sentences can have different interpretations in different contexts. There are puns, jokes, and metaphors as well as grammar with all its exceptions. Raina and Krishnamurthy (2022) have highlighted a few prevailing tasks NLP is used for: sentiment analysis, document classification, and intent classification. Sentiment analysis refers to the overall sentiment of the text i.e., does the text convey positive or negative feelings. Document classification consists of how the document is classified and categorised. Intent classification assigns the intent of the text. Text prediction and autocomplete are also considered major challenges in the field of NLP.

The NLP used these days is usually statistical NLP (IBM 2020). This is a combination of deep learning and machine learning where elements of text are automatically extracted, classified, and labelled. Each element is then assigned a statistical likelihood of a possible meaning. This enables NLP to learn to become more accurate in its predictions, while

analysing more and more datasets. NLP is usually thought using deep learning methods which in turn are thought using artificial neural networks. In this particular case, deep learning models are used to train a model that can represent words as vectors (for more details on this process see section 3.2) (Yang et al. 2022). These vectors are then passed to the NLP algorithm for processing.

Deep learning is done using neural networks. A neural network is a network structure that mimics the human brains neural networks. It is comprised of multiple connected units called artificial neurons. Based on the way the neurons are connected to each other, the neural network configuration is determined. A neural network has three layers of neurons: input layers, hidden layers, and output layers. For instance, a deep neural network is a network that has two or more hidden layers (Bayraci & Susuz 2019). Each neuron performs a calculation, which is passed on to other neurons deeper in the network. A neuron can be considered as a linear function. This is shown in equation 1. Input signals are assigned weights and added a constant, bias, (i.e., input times its weight plus bias) and adjusting these weights and biases, is how models are trained.

$$y = xw + b \quad (1)$$

where x is the input, w is the weight assigned to it, b is the bias and y is the output. These y output signals are then summed up and passed to an activation function. In this case the activation function is used to map out a probability distribution of whether an article is fake or real. An illustration of the calculations in a neuron are showed in figure 1.




Figure 1. Illustration of a neuron as a linear function (Srijayjk 2020).

While this is just a basic artificial neural network, there are many different types and models which also vary based on the usage.

2.3 Related works

While there are many works done on the specific models and their architectures for fake news detection, there are not so many works describing common methods and techniques that are generally needed for good models. Medeiros and Braga (2020) did a systematic review of fake news detection in social media. According to their research most studies were based on Twitter and the Chinese social media platform Weibo, and the best results were achieved using an LSTM (Long Short-Term-Memory) based classifier. Saquete et al. (2019) also conducted a systematic review of state-of-the-art classifier methods as well as subtasks like feature extraction methods.

Saha et al. (2022) did a comparative analysis of fake news methods. They compared deep learning and traditional machine learning methods as well as these methods together and used accuracy as the evaluation metric. According to their research deep learning methods performed better than traditional machine learning methods but the best accuracy was obtained from using these methods together. A benchmark study of machine learning models for online fake news detection created by Khan et al. (2021) found that the BERT based models were the best option.

Most datasets have articles in English or Chinese, so most models are only trained with these languages. Faustini and Covões (2020) experimented on five datasets in three different languages: Portuguese, English and Bulgarian. For classifiers, they used traditional machine learning methods. Besides classifiers, they also experimented with different feature extractors. Hämäläinen and Alnajjar (2021) reviewed the current state of Finnish NLP. They gathered studies from different NLP tasks.

3 Model techniques

In this section, the technicalities and most typical steps described in most studies are presented and explained in more detail. This is done using a thematical analysis and researching papers that have proposed a deep learning model for detecting fake news. Common steps include gathering the dataset, pre-processing the data, feature extraction, training the classifier and producing the end result. This is visualized in figure 2.

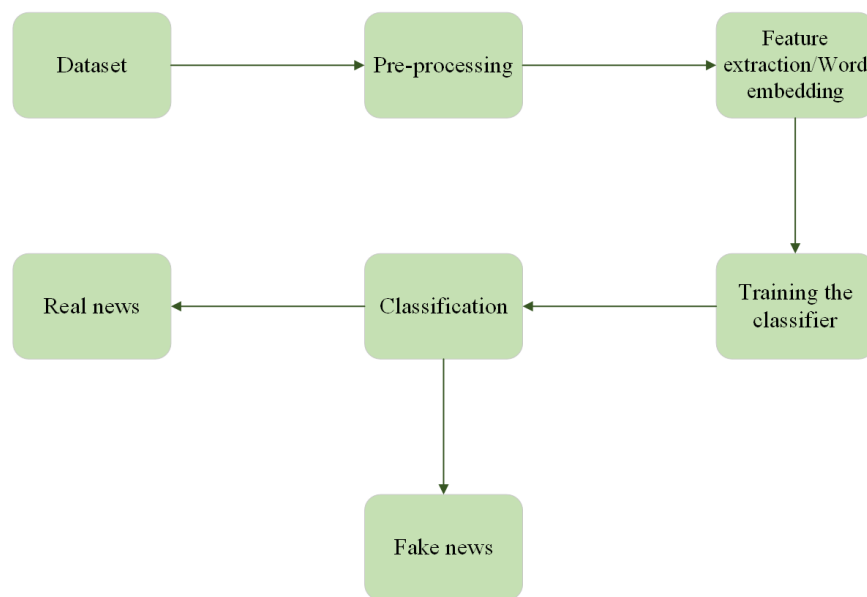


Figure 2. Visualization of data processing in detection methods.

3.1 Dataset and pre-processing

Every proposed model has to be trained first with a dataset that contains fake and real news. The dataset can either include both real and fake news or they can be two separate sets. Datasets are split into two sets, one for training the model and the other for testing. These datasets contain on minimum the article's title, body, and a label of either fake or real. The testing dataset does not include a pre-set label, but other attributes are the same. Kaggle.com is a website that has multiple datasets for this purpose, and these are used in most related works. Additionally, some studies have then gathered their own dataset of real news from the sites such as New York Times, the Guardian, and Reuters among others. (Ahmed, Traore

& Saad 2017b; O'Brien et al. 2018; Yang et al. 2022; Ahmed, Traore & Saad 2017a; Khan et al. 2021.) For the deep learning methods to train properly, it is essential that a dataset is considerably large. For instance, Chauhan & Palivela (2021) used a dataset that comprised of 40 000 articles in total which had approximately 20 000 fake articles and 20 000 real articles. The rough medium seems to be a dataset of 20 000 articles in total.

After gathering the required dataset, the data needs to be processed into a format that the machine can then use. Pre-processing also guarantees more accurate results. Pre-processing usually includes removing stop words, lowering the case, punctual removing, and sentence segmentation. Stop words are insignificant words that are used to help connect thought or to assist in the sentence structure (Ahmed, Traore & Saad 2017b). Examples of stop words are a, an, the, about, by, for, from, too, what, when etc. Text is separated into individual words after which, a popular method is using the Python natural language toolkit that contains a list of stop words which the text can be referenced against to remove words that are in this list (Chauhan & Palivela 2021; Alsuliman et al. 2022). Text's case is lowered, converting numbers into written numbers, and other non-letter characters are removed. When these steps are finished lemmatization or stemming is conducted. Stemming means that a word is stripped of its suffixes and prefixes, resulting in the original and most basic form of the word. For instance, sitting and sat would be reduced to sit. (Ahmed, Traore & Saad 2017b). In comparison, lemmatization takes into consideration the morphological analysis of each word and returns the base or dictionary form of the word (Balakrishnan & Lloyd-Yemoh 2014). These essentially produce the same result but, in some cases, lemmatization has proven to be more accurate according to Balakrishnan and Lloyd-Yemoh (2014). Now the texts are ready to be vectorized.

3.2 Word embedding and feature extraction

To start feature extraction, the text data needs to be tokenized. "Tokenization is the process of representing each word in the form of a number." (Chauhan & Palivela 2021). This still leaves a dataset of monumental size that leads to a high computational burden for the learning process (Ahmed, Traore & Saad 2017b). To overcome this issue feature extraction is the required next step. Popular feature extraction methods are Term Frequency (TF) and Term Frequency-Inverted Document Frequency (TF-IDF). TF uses the word counts in the

texts to determine how similar they are to one another. Depending on the number of times a word appears in the text, the word is given a weight. This is done using the following equation:

$$TF(t, d) = \frac{\text{Term } t \text{ frequency in document } d}{\text{Total words in document } d} \quad (2)$$

(Alsuliman et al. 2022). TF-IDF is used to measure how important a term is in the document. IDF scales up the rare words and weights down the frequent ones. To do this IDF focuses on how many documents the word appears in. IDF is measured using the following equation:

$$IDF = \text{Log} \frac{\text{Total documents}}{\text{Total documents with term } t+1} \quad (3)$$

TF-IDF is therefore TF multiplied by IDF:

$$TF - IDF = TF(t, d) * IDF \quad (4)$$

(Alsuliman et al. 2022). If a word does not appear in a document, it is denoted as 0.

Word embedding acts as a feature extraction method. Word embedding is the process of vectorizing words and creating a vector space where words with similar definitions are positioned close to each other in the vector space (Goldani, Momtazi & Safabakhsh 2021). Some examples of word embedding algorithms include gloVe (Global Vectors for Word Representation) and word2vec. Word2vec is a group of models created by training neural networks that learn word embeddings from a large set of raw text. This then creates a vector space with several hundred dimensions and appoints each word a unique vector. GloVe on the other hand is trained on a global word to word co-occurrence statistics. Meaning that while word2vec is trained on local statistics, i.e., local context, gloVe incorporates world-wide statistics, i.e., wider word co-occurrences. (Goldani, Momtazi & Safabakhsh 2021; Chauhan & Palivela 2021; Pennington, Socher & Manning 2014).

Another option for data segmentation is the use of n-grams (Li, Yu & Zhang 2022). N-grams are words or phrases that are co-occurring or continuous sequences of n items. The n in n-gram defines the length of these phrases. For instance, if N = 1, also called a unigram, unique words are taken into consideration. If N = 2, also called a bigram, two-word combinations are considered. Figure 3. visualises n-grams.

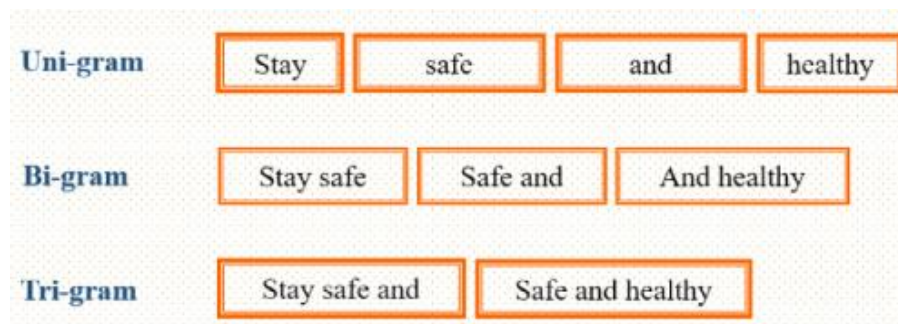


Figure 3. Example of n-grams (Chauhan & Palivela 2021).

An n-gram model predicts the most likely adjacent word to the examined item. N-grams can be scaled as large as necessary.

In conclusion, some type of feature extraction needs to be done before continuing to the next step, classification. Usage of either TF-IDF or some form of word embedding seems to be distributed quite evenly in related works. Preference between word2vec and gloVe seems also to be quite even, although some have suggested that gloVe is more accurate (Chauhan & Palivela 2021; Pennington, Socher & Manning 2014).

3.3 Classification methods

Classification is an NLP task where machine learning algorithms use pre-categorized training datasets to classify future datasets. Algorithms predict in what predetermined category the data will fall into. There are many proposed deep learning-based methods and models, so the ones mentioned here are the state-of-the-art and the models that produced the highest accuracy scores. While there are also traditional machine learning models like linear support vector machine, naive Bayes, decision tree and k-nearest neighbours these have not performed as well compared with deep learning-based methods (Khan et al. 2021).

3.3.1 CNN-based

CNN (convolutional neural network) is typically used in image recognition and computer vision tasks, but it has proven to be useful in text classification tasks as well. CNN is a type of DNN, and it has a separate convolutional layer. When using CNN for text classification

the input layer is the word embeddings, i.e., the word vector matrix. Feature extraction is done in the convolution layer by applying filters over the input data. This data is then also pooled to reduce the dimensional complexity and lessening the computational burden for the learning process. Normally max pooling is done. Max pooling selects the maximum value from the vector that represents the most used word in the text. These maximum values are then put through an activation function to produce the output. (Choubey 2020). A simple CNN architecture is presented in figure 4.




Figure 4. Visualization of how CNN works.

O'Brien et al. (2018) studied the transferability of the learning process to novel news topics. A model is initially trained on a dataset of specific topics, after which new topics are introduced to the model. They then studied how the learning process of the model was altered because of these new topics. They use CNN as the classifier and word2vec as word embedding in their proposed detection model. Their model uses only language to detect fake news (no syntax, semantics, or source analysis) and thus proves that CNN performs well when introducing new topics.

Yang et al. (2022) created a CNN based classification model they named TI-CNN (Text and Image information based Convolutional Neural Network). They consider explicit and latent features as well as text and image features in their model. Latent feature, i.e., hidden feature, is like the common topic in a bunch of words. For instance, if you have car, bus, tram, train, the latent feature would be vehicle. They use two parallel CNN layers to extract the latent features from texts and images. Explicit and latent feature spaces are projected together after which text and image feature spaces are then fused together for the fake news detection. It

is stated that image-only information is insufficient in detection, but text and image together work better than some other methods they compared TI-CNN with.

3.3.2 LSTM-based

LSTM (Long Short-Term Memory) is a type of recurrent neural network (RNN) but instead of neural network layers it has LSTM cell blocks: the input gate, the forget gate and the output gate. Unlike DNN, RNNs have a feedback loop meaning that it learns based on the previous information. RNN takes the last output and concatenates it with the new input. LSTM is a sequential model meaning that it can process entire sequences of data (Chauhan & Palivela 2021). The concatenation of new and old input data flows through the input gate to the forget gate where it is determined if some parts of the data are kept or lost. Data then flows to the output gate and determines what parts are actually allowed as an output. This is done with a series of sigmoid and tanh functions as well as multiplication and addition operations. In figure 5. an LSTM cell is illustrated. Memory cell input refers to the old output.

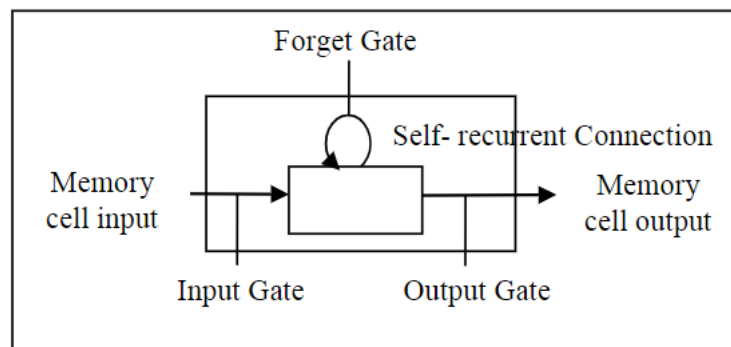


Figure 5. An LSTM cell (Bahad, Saxena & Kamal 2019).

Chauhan and Palivela (2021) created a stacked LSTM architecture which has two LSTM layers. They used gloVe for word embedding and n-grams to enhance the tokenization. Their trained model received an accuracy score of 99.88 % which is the highest accuracy percentage according to their comparative analysis with various research works.

Bahad, Saxena & Kamal (2019) use a bi-directional LSTM (bi-LSTM) in their fake news detection model. Bi-LSTM is two RNNs working simultaneously, one going forward (past

to future) and the second going backward (future to past) so that the network has both information at the same time. Their model received an accuracy of 91.08 % and 98.75 % on the two datasets used. They also trained CNN, vanilla RNN and unidirectional LSTM to compare their bi-LSTM against. The comparative results show that unidirectional and bi-LSTM had very similar results and that unidirectional LSTM was more accurate in the first dataset and bi-LSTM was more accurate in the second dataset.

3.3.3 CSI

According to Ruchansky, Seo and Liu (2017) three general characteristics of fake news are text, response, and source. Based on these they built a hybrid deep model, CSI. It is composed of three modules: Capture, Score and Integrate. First module focuses on the response and text. This is done using LSTM to capture the pattern of user activity on a given article. The second module focuses on the source based on users' behaviour and activity, giving scores to users based on the content they engage with and suspiciousness. The third module is the two modules integrated together to produce a prediction of either fake or real. Advantages of CSI is that it requires fewer parameters and less training. Accuracy scores were 95.3 % on the Chinese social media platform Weibo and 89.2 % on Twitter. (Ruchansky, Seo & Liu 2017).

3.3.4 BERT-based

BERT (Bidirectional Encoder Representations from Transformers) is a pre-trained word embedding model for NLP created by Google. While other models like CNN, RNN and LSTM have many different versions for different tasks, BERT has a unified architecture and only needs fine-tuning across varied tasks. As opposed to directional models, BERT does not read the text input left-to-right or right-to-left but instead the whole sequence at once. This lets it learn the context of a word based on all surroundings, both left and right of the word. BERT uses masked language model to randomly mask some input tokens and then predicts solely based on the context what words the tokens were. BERT also uses next sentence prediction task to train text-pair representations. Masked language model and next

sentence prediction are trained together. While this creates more accurate predictions it also requires more training time. (Devlin et al. 2019).

Deng and Wang (2022) combined the CSI framework and a BERT based encoder (CSIBERT) in their model and outperformed other models including Ruchansky's CSI model. As opposed to the CSI model they used a bi-LSTM layer for the input of the capture module. Their accuracy score was 97.1 %. Another BERT based model is by Kaliyar, Goswami and Narang (2021). They combined BERT and three CNN layers to make the FakeBERT model which was 98.9 % accurate.

4 Comparison of results

Determining what model and architecture to use for fake news detection is not so unequivocal. Since there are so many possibilities of machine learning techniques for classifiers, different feature extraction methods as well as many datasets with slightly different data, it is impossible to say which combination is the best. Nonetheless, in table 1. is gathered some proposed models with the highest accuracy scores from the aforementioned deep learning-based classifiers. An accuracy score is mentioned in most papers and is possibly the best evaluation metric to compare models against each other.

Table 1. Highest accuracy scores.

Deep learning based	Creators	Dataset size	Accuracy
Stacked LSTM	Chauhan and Palivela (2021)	40 000	99.88 %
Bi-LSTM	Bahad, Saxena and Kamal (2019)	6 200 and 3 300	98.75 %
FakeBERT	Kaliyar, Goswami and Narang (2021)	20 800	98.9 %
CSIBERT	Deng and Wang (2022)	4 600	97.1 %
CSI	Ruchansky, Seo and Liu (2017)	4 600	95.3 %
CNN	O'Brien et al. (2017)	24 000	93.5 %
TI-CNN	Yang et al. (2022)	20 000	92.2 %

There are four scores that need to be considered. True Positives (TP) mean that the actual class and the predicted class are both true. True Negatives (TN) mean that the actual class and the predicted class are both false. False Positives (FP) mean that the actual class is not the same as the expected class. And lastly, False Negatives (FN) mean that the actual class is true, but the expected class is false. Accuracy is then measured according to equation 5.

$$Accuracy = \frac{(TP+TN)}{(TP+FP+TN+FN)} \quad (5)$$

(Khan et al. 2022). Accuracy score is also the only common metric all these works share. The LSTM based models have the highest accuracy followed by BERT, CSI and CNN based. As seen in table 1. the stacked LSTM model proposed by Chauhan and Palivela has the highest accuracy score of 99.88 % as well as the largest training dataset of 40 000 articles.

The accuracy of a model can be affected by for instance the size of the dataset. The bi-LSTM model created by Bahad, Saxena and Kamal (2019) was trained on two datasets separately which were the size of about 6 000 and 3 000 articles. The test run on the dataset of 3 000 values received the acclaimed accuracy of 98.75 % and the test run on the bigger dataset was 91.08 % accurate. In their comparative tests, the unidirectional LSTM performed even better with an accuracy of 91.48 %. Likewise, the CSI and CSIBERT classifiers were trained using a dataset gathered from the Chinese social media platform Weibo. This contains 4 600 articles. Teaching the network with a smaller dataset means that the networks knowledge becomes more limited and one-sided. It is also easier to get good results when the network does not have so much material to work with. This leads the model to possibly not being as useful for everyday use. Nevertheless, this model is state-of-the-art for bi-LSTM-based fake news detectors for its proposed classifier architecture.

The datasets and their sizes vary quite a lot. CNN and TI-CNN used the same dataset for fake news from Kaggle and then gathered their own real news from various reputable news websites like the New York Times and Washington Post. CSI and CSIBERT also used the same dataset, but all other models were trained on different datasets. The respectable sizes can be observed in table 1.

Models use different feature extraction methods thus a specific method cannot be pinpointed as a factor for a successful model. Stacked LSTM and bi-LSTM use gloVe. The BERT based models use a BERT base encoder as the text feature extractor. CSI uses doc2vec, the CNN model and TI-CNN uses word2vec. Doc2vec creates a feature vector for the entire document while word2vec creates it for specific words. Further research is needed for optimising the classification models and different feature extraction methods.

Seeing as the stacked LSTM model has the highest accuracy score and it was trained over the largest dataset it shows a lot of promise for public use. The end goal would be a model where anyone can input any article to the system, and it would output if the content is fake or real.

4.1 Finnish fake news detection

There are not many fake news detection models for Finnish. The closest one to fake news detection is by Hämäläinen et al. (2021). This is a rumour detection model that uses FinBERT, a BERT model trained on various texts in Finnish which can be fine-tuned for different NLP tasks. Their experiments show that the multilingual BERT which was trained over Finnish among others achieved the best accuracy of 97.2 % for factual texts. For rumour detection FinBERT outperformed other models with 96 % accuracy. This is explained by the training datasets used for these models. FinBERT was trained over Suomi24 which is a forum where anyone can post almost anything, and multilingual BERT was trained over Wikipedia pages. Both models were compared with LSTM and LSTM + word2vec methods.

For more research to be done on fake news detection in Finnish using different models, a dataset of Finnish fake and real news is needed. Hämäläinen et al. (2021) published the dataset they used in their study, but this contains rumours not fake news per se. For data pre-processing, there is a lemmatization toolkit in Python for Finnish (Antupis Blog 2019) and tokenization can be done using the Python natural language tool kits. A blog post published by Kanstren (2017) talked about his experiments with word2vec for Finnish which yielded satisfactory results. Hämäläinen et al. also used word2vec when training the LSTM classifier.

A model is definitely possible to create for Finnish fake news detection. It would require creating a new dataset but otherwise all other tools already exist. The only hindering factor is the amount of interest researchers in the Finnish NLP community have for this subject. While there is fake news even in Finland, the threat is not so concerning to society as it is in bigger countries. Finland has a national initiative called responsible journalism (vastuullista journalismia) where well written and source critical articles are marked with a little blue tick to indicate that the article is reliable and truthfully written. Source criticism is also taught since elementary school and all throughout the education system. Meaning that the need for a Finnish fake news detection system is not on the top of NLP researchers priority list. Although, during the COVID-19 pandemic the anti-vaccine movements were spreading here as well, so this might need to change.

5 Discussion and conclusions

Fake news has become a threat to society. The surge of fake news started during the 2016 US presidential elections, re-enforced during the COVID-19 pandemic, and is still wreaking havoc today. To prevent the spread of fake news an automated detection system is needed. While there are already manual fake news detection websites like PolitiFact, an automated detector would scale larger, and the results could be available instantly. Fake news can be detected using NLP techniques. The NLP task, text classification is the best and most used method. A classifier needs to be trained over a large dataset of pre-processed data. Data processing also includes vectorization of words and creating a vector space where words with similar meanings are close together, this is also called feature extraction. The models with the highest accuracy used Long Short-Term Memory (LSTM), Bidirectional Encoder Representations from Transformers (BERT) and Convolutional Neural Network (CNN) based classifiers. A model proposed by Chauhan and Palivela used two parallel layers of LSTM creating a stacked LSTM, that has an accuracy of 99.88 % which is the highest accuracy rate compared to all other models. The stacked LSTM model was also trained over a dataset of 40 000 articles, meaning that the model probably performs well on articles outside the dataset as well. This is a promising model for eventual public use on all articles.

For the comparison of models' results, accuracy score was picked as the defining metric because this was the only one in common that all models were presented with. There is also a few other metrics that give more detailed information on the model's progress like precision, recall and F1. Precision calculates the accuracy of positive predictions, meaning how accurate were all the predictions of real news. Recall measures the ratio of correctly predicted instances and all instances of a given class. So, for real news the given class would be the correctly predicted actually real news (TP) and the falsely predicted but actually real news (FN). For fake news it is the opposites i.e., TN and FN. F1-score is the mean of the precision and recall. (Khan et al. 2021). This brings up the dilemma of which is the lesser evil: classifying real news as fake or fake news as real. Classifying fake news as real can be considered more harmful since the reader would then base their knowledge of that subject as real on other fake news as well. Classifying real news as fake is the more cautious choice.

Out of the two psychological factors presented by Shu et al. (2017) one is confirmation bias. This means that people prefer information that confirms their bias giving them a positive emotional feedback where they get to feel that they are in the right and everyone else is in the wrong. This also means that correcting them and saying that they are in the wrong leads to them rejecting the ideas and causing more harm than good. Nevertheless, this only applies to the devout believers and correcting and warning all users of fake news prevents new people to start following fake news agendas.

Another remark is that models could be biased since most datasets include political news. Fake news is generally spread for Trump and Republican supporters and the real news selected for the models were from websites that are either unbiased or inclined to favour Democrats. This could lead the classifiers to be biased toward Democrats and the more liberal and left wing of the political scene. Training the models over a dataset with more varied topics or from both political sides would prevent this.

5.1 Central conclusions to research questions

In the beginning of this thesis three research questions were set which now have been answered.

- How is NLP used to detect fake news online?

Text classification is an NLP task which is utilised in fake news detection. There are more options like entity recognition, but the models reviewed in this work used text classification. Neural networks are built as classifiers which are taught over datasets of fake and real news. This data in the datasets needs to be cleaned and pre-processed first after which the text is vectorized based on a feature space. Classifiers then produce an output of a prediction if an article is fake or real based on their learning.

- What types of NLP have been used for this purpose and what have been the results?

There are many options for how a classifier is built but the ones producing the best results were LSTM, BERT, CSI, and CNN based models. LSTM based models performed the best with the highest accuracy score of 99.88 % on the stacked LSTM model. Other models performed also well, the lowest accuracy score being 92.2 % on the TI-CNN model.

- Can these be used for Finnish language platforms?

Definitely. There already exists a rumour detection model that uses FinBERT and was 96 % accurate. A fake news detector does not seem to exist yet for Finnish language but there are tools for the data processing and implementing for instance the stacked LSTM classifier is possible. A dataset of Finnish fake and real news does not exist so this would have to be gathered.

5.2 Research bias and reliability

Initial research boundary limited the research to state-of-the-art methods, otherwise the research is unbiased. The exemplar classifiers were picked due to their newness, best results and because those works were cited in many other works as well. The methods for feature extractors and classifiers i.e., gloVe, CNN, LSTM etc., were picked due to them being mentioned the most frequently in all works considering the topic. Reliability is based on source criticism.

5.3 Future outlook and prospects

It is naïve to say that fake news will be completely eradicated in the future. But since fake news detection is researched and different types of solutions are proposed frequently, it is probable that stopping and warning about fake news is more sufficient and done in a larger scale. An example of fake news detection and prevention could be seen during the COVID-19 pandemic when Twitter added labels to tweets that had misleading information about the virus. There are already Chrome extensions for fake news detection as well, like BRENDA which gives the user feedback on the article they are reading (Botnevik, Sakariassen & Setty 2020). Using these solutions in a larger scale and implementing fake news detection methods in all social medias as well as educating people on source criticism could diminish the effect fake news has on the public today.

5.4 Future works

For future works, a wider scope of research papers could be examined as well as looking at more classifier methods. In addition, a technical experiment could be conducted based on the best models for Finnish.

References

- Ahmed, H., Traore, I., Saad, S., 2017a. ‘Detection of Online Fake News Using N-Gram Analysis and Machine Learning Techniques’. *International Conference on Intelligent, Secure, and Dependable Systems in Distributed and Cloud Environments (ISDDC 2017)*. 11 October. p. 127–138. [Cited 2022-11-07] DOI: https://doi.org/10.1007/978-3-319-69155-8_9
- Ahmed, H., Traore, I., Saad, S., 2017b. ‘Detecting opinion spams and fake news using text classification’. *Security and Privacy*, 1, e9, [Online]. [Cited 2022-11-07] DOI: <https://doi.org/10.1002/spy2.9>
- Alsuliman, F. et al. 2022. ‘Social Media vs. News Platforms: A Cross-analysis for Fake News Detection Using Web Scraping and NLP’. *Proceedings of the 15th International Conference on Pervasive Technologies Related to Assistive Environments (PETRA 2022)*. 29 June. p. 190–196. [Cited 2022-11-07] DOI: <https://doi.org/10.1145/3529190.3534755>
- Antupis, 2019. ‘Finnish Lemmatization With Python’, *Antupis Blog*, 12th June. [Blog]. [Cited 2022-11-24]. Available:
- Bahad, P., Saxena, P., Kamal, R., 2019. ‘Fake News Detection using Bi-directional LSTM-Recurrent Neural Network’. *Procedia Computer Science*, 165, p. 74–82, [Online]. [Cited 2022-11-15] DOI: <https://doi.org/10.1016/j.procs.2020.01.072>
- Balakrishnan, V., Lloyd-Yemoh, E., 2014. ‘Stemming and Lemmatization: A Comparison of Retrieval Performances’. *Lecture Notes on Software Engineering* 2(3), p. 262–267. [Cited 2022-11-08] DOI: <https://doi.org/10.7763/LNSE.2014.V2.134>
- Bayraci, S., Susuz, O., 2019. ‘A Deep Neural Network (DNN) based classification model in application to loan default prediction’. *Theoretical and Applied Economics*, 26(4(621)), p. 75–84. [Cited 2022-11-12]. Available: [https://ideas.repec.org/a/agr/journal/vxxvvi2019i4\(621\)p75-84.html](https://ideas.repec.org/a/agr/journal/vxxvvi2019i4(621)p75-84.html)
- Bessi, A., Ferrara, E., 2016. ‘Social Bots Distort the 2016 US Presidential Election Online Discussion’. *First Monday*, 21(11-7), [Online]. [Cited 2022-10-26]. Available: <https://firstmonday.org/article/view/7090/5653>

Botnevik, B., Sakariassen, E., Setty, V., 2020. 'BRENDA: Browser Extension for Fake News Detection'. *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2020)*. 25 July. p. 2117–2120. [Cited 2022-12-01] DOI: <https://doi.org/10.1145/3397271.3401396>

Chauhan, T., Palivela, H., 2021. 'Optimization and improvement of fake news detection using deep learning approaches for societal benefit'. *International Journal of Information Management Data Insights*, 1(2), 100051, [Online]. [Cited 2022-11-07] DOI: <https://doi.org/10.1016/j.jjime.2021.100051>

Choubey, V., 2020. 'Text classification using CNN'. *Medium*, 8th July [blog]. [Cited 2022-11-13]. Available: <https://medium.com/voice-tech-podcast/text-classification-using-cnn-9ade8155dfb9>

Deng, Y., Wang, S.-W., 2022. 'Detecting Fake News on Social Media by CSIBERT'. *Proceedings of the 2022 6th International Conference on Deep Learning Technologies, (ICDLT 2022)*. 26–28 July. p. 142–148. [Cited 2022-11-20] DOI: <https://doi.org/10.1145/3556677.3556698>

Devlin, J. et al. 2019. 'BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding'. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. June 2019. p. 4171–4186. [Cited 2022-11-20] DOI: <https://doi.org/10.18653/v1/N19-1423>

Edkins, B., 2016. 'Americans Believe They Can Detect Fake News. Studies Show They Can't'. *Forbes*, 20 December, [article]. [Cited 2022-10-24]. Available: <https://www.forbes.com/sites/brettedkins/2016/12/20/americans-believe-they-can-detect-fake-news-studies-show-they-cant/>

Faustini, P.H.A., Covões, T.F., 2020. 'Fake news detection in multiple platforms and languages'. *Expert Systems with Applications*, 158, 113503, [Online]. [Cited 2022-11-21] DOI: <https://doi.org/10.1016/j.eswa.2020.113503>

Ferrara, E. et al. 2016. 'The rise of social bots'. *Communications of the ACM*, 59(7), p. 96–104, [Online]. [Cited 2022-10-26] DOI: <https://doi.org/10.1145/2818717>

Goldani, M.H., Momtazi, S., Safabakhsh, R., 2021. 'Detecting fake news with capsule neural networks'. *Applied Soft Computing*, 101, 106991, [Online]. [Cited 2022-11-09] DOI: <https://doi.org/10.1016/j.asoc.2020.106991>

Hämäläinen, M. et al. 2021. 'Never guess what I heard... Rumor Detection in Finnish News: a Dataset and a Baseline'. *Proceedings of the Fourth Workshop on NLP for Internet Freedom: Censorship, Disinformation, and Propaganda*. June 2021. p. 39-44. [Cited 2022-11-24] DOI: <https://doi.org/10.18653/v1/2021.nlp4if-1.6>

Hämäläinen, M., Alnajjar, K., 2021. 'The Current State of Finnish NLP'. *Proceedings of the Seventh International Workshop on Computational Linguistics of Uralic Languages*. September 2021, p. 54–61. [Cited 2022-11-21] DOI: https://doi.org/10.26615/978-1-954085-82-4_007

<https://antupis.github.io/lemmatization/finnish/2019/06/12/Lemmatizing-finnish-text.html>

IBM Cloud Education, 2020. 'What is Natural Language Processing?' [website]. [Cited 2022-10-28]. Available: <https://www.ibm.com/cloud/learn/natural-language-processing>

Kaliyar, R.K., Goswami, A., Narang, P., 2021. 'FakeBERT: Fake news detection in social media with a BERT-based deep learning approach'. *Multimed Tools Applications*, 80, p. 11765–11788, [Online]. [Cited 2022-11-20] DOI: <https://doi.org/10.1007/s11042-020-10183-2>

Kanstren, T., 2017. 'Word2Vec with some Finnish NLP'. *Random experiments in software engineering*, 19th November, [blog]. [Cited 2022-11-24]. Available: <https://teemukanstren.com/2017/11/19/word2vec-with-some-finnish-nlp/>

Khan, J.Y. et al. 2021. 'A benchmark study of machine learning models for online fake news detection'. *Machine Learning with Applications*, 4, 100032, [Online]. [Cited 2022-11-07] DOI: <https://doi.org/10.1016/j.mlwa.2021.100032>

Li, T., Yu, J., Zhang, H., 2022. 'Web of things based social media fake news classification with feature extraction using pre-trained convoluted recurrent network with deep fuzzy learning'. *Theoretical Computer Science*, 931, p. 65–77, [Online]. [Cited 2022-11-09] DOI: <https://doi.org/10.1016/j.tcs.2022.07.031>

Medeiros, F.D.C., Braga, R.B., 2020. 'Fake News Detection in Social Media: A Systematic Review'. *XVI Brazilian Symposium on Information Systems, (SBSI 2020)*. 03 November. p. 1–8. [Cited 2022-11-21] DOI: <https://doi.org/10.1145/3411564.3411648>

O'Brien, N. et al. 2018. 'The Language of Fake News: Opening the Black-Box of Deep Learning Based Detectors'. *32nd Conference on Neural Information Processing Systems, (NIPS 2018)*. 2–8 December. [Cited 2022-11-07]. Available: <https://dspace.mit.edu/handle/1721.1/120056>

Pennington, J., Socher, R., Manning, C., 2014. 'GloVe: Global Vectors for Word Representation'. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*. 26–28 October. p. 1532–1543. [Cited 2022-11-10] DOI: <https://doi.org/10.3115/v1/D14-1162>

Raina, V., Krishnamurthy, S., 2022. 'Natural Language Processing'. *Building an Effective Data Science Practice*, 1, p. 63–73, [Online]. [Cited 2022-10-28] DOI: https://doi.org/10.1007/978-1-4842-7419-4_6

Ruchansky, N., Seo, S., Liu, Y., 2017. 'CSI: A Hybrid Deep Model for Fake News Detection'. *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, (CIKM 2017)*. 6–10 November. p. 797–806. [Cited 2022-10-25] DOI: <https://doi.org/10.1145/3132847.3132877>

Saha, A. et al. 2022. 'A Comparative Analysis on Fake News Detection Methods'. *Proceedings of the 2nd International Conference on Computing Advancements, (ICCA 2022)*. 10–12 March. p. 392–399. [Cited 2022-11-21] DOI: <https://doi.org/10.1145/3542954.3543010>

Saquete, E. et al. 2020. 'Fighting post-truth using natural language processing: A review and open challenges'. *Expert Systems with Applications*, 141, 112943, [Online]. [Cited 2022-11-21] DOI: <https://doi.org/10.1016/j.eswa.2019.112943>

Shu, K. et al. 2017. 'Fake News Detection on Social Media: A Data Mining Perspective'. *ACM SIGKDD Explorations Newsletter*, 19(1), p. 22–36, [Online]. [Cited 2022-10-24] DOI: <https://doi.org/10.1145/3137597.3137600>

Srijay, K., 2020. 'Activation Functions in DNN'. *Medium*, 14th June. [Cited 2022-11-12]. Available: <https://srijayjk.medium.com/activation-functions-in-dnn-21729d529364>

Watson, C.A., 2018. 'Information Literacy in a Fake/False News World: An Overview of the Characteristics of Fake News and its Historical Development'. *International Journal of Legal Information*, 46(2), p. 93–96, [Online]. [Cited 2022-10-24] DOI: <https://doi.org/10.1017/jli.2018.25>

Yang, Y. et al. 2022. 'TI-CNN: Convolutional Neural Networks for Fake News Detection'. [Online]. [Cited 2022-10-25] DOI: <https://doi.org/10.48550/arXiv.1806.00749>

Ye, J., Skiena, S., 2019. 'MediaRank: Computational Ranking of Online News Sources'. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, (KDD 2019)*. 4–8 August. p. 2469–2477. [Cited 2022-10-27] DOI: <https://doi.org/10.1145/3292500.3330709>

Zhou, X., Zafarani, R., 2020. 'A Survey of Fake News: Fundamental Theories, Detection Methods, and Opportunities'. *ACM Computing Surveys*, 53(5), p. 1–40, [Online]. [Cited 2022-10-24] DOI: <https://doi.org/10.1145/3395046>