

Article

Automatic Fake News Detection for Romanian Online News

Marius Cristian Buzea ¹, Stefan Trausan-Matu ^{1,2,*}  and Traian Rebedea ¹ 

¹ Department of Computer Science and Engineering, University Politehnica of Bucharest, 060042 Bucharest, Romania; bumarius@gmail.com (M.C.B.); traian.rebedea@upb.ro (T.R.)

² Research Institute for Artificial Intelligence “Mihai Drăganescu” of the Romanian Academy, 050711 Bucharest, Romania

* Correspondence: stefan.trausan@upb.ro or trausan@gmail.com

Abstract: This paper proposes a supervised machine learning system to detect fake news in online sources published in Romanian. Additionally, this work presents a comparison of the obtained results by using recurrent neural networks based on long short-term memory and gated recurrent unit cells, a convolutional neural network, and a Bidirectional Encoder Representations from Transformers (BERT) model, namely RoBERT, a pre-trained Romanian BERT model. The deep learning architectures are compared with the results achieved by two classical classification algorithms: Naïve Bayes and Support Vector Machine. The proposed approach is based on a Romanian news corpus containing 25,841 true news items and 13,064 fake news items. The best result is over 98.20%, achieved by the convolutional neural network, which outperforms the standard classification methods and the BERT models. Moreover, based on irony detection and sentiment analysis systems, additional details are revealed about the irony phenomenon and sentiment analysis field which are used to tackle fake news challenges.

Keywords: fake news; convolutional neural network; machine learning; deep learning; RoBERT



Citation: Buzea, M.C.; Trausan-Matu, S.; Rebedea, T. Automatic Fake News Detection for Romanian Online News. *Information* **2022**, *13*, 151. <https://doi.org/10.3390/info13030151>

Academic Editor: Kostas Stefanidis

Received: 18 February 2022

Accepted: 10 March 2022

Published: 14 March 2022

Publisher’s Note: MDPI stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Copyright: © 2022 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Over the recent years, artificial intelligence (AI) brought important changes in the domain of information technologies and architectures, such as using and developing intelligent transportation systems, virtual personal assistants, robotic surgery, and maybe with the greatest impact on our lives, natural language processing [1].

Nowadays, due to the internet, the quality and quantity of the news increases every day, and the way that the consumer accesses and manages daily online information is constantly changing. Young people, especially the inexperienced ones, use social media platforms, mobile applications, or simple and dynamic websites to extract the necessary information quickly and easily, many times without discernment. The diversity of the online news may increase the engagement in democratic elections, giving everyone the opportunity to get involved or even change opinions and mentalities. However, new technologies and features can be used through social media platforms to spread fake news on a large scale, creating personalized information and becoming more effective in misinformation campaigns. Therefore, credible and reliable sources of information are needed so that the public does not fall prey to the intentions of those interested in manipulating reality. Some researchers [2] have also suggested that populist politicians use fake news in order to undermine authority. For example, recent research suggests that fake news is used to doubt some sources of information that in the past had been considered trusted in several fields, such as the scientific community or journalism [3].

Fake news can function as propaganda or misinformation, but it always appeals to the emotions of the public and the intent to cover rational responses, analysis, and comparison of information from several sources, encouraging inflammation and outrage and can easily lead to conspiracy theories and partisan biased content that negatively affects

social security. Hence, to tackle fake news, the researchers involve several approaches such as text classification, network analysis, or reviews and evaluations [4]. Moreover, Caplan et al. [5] described how companies and AI researchers should define fake news by type. There are other ways to address misinformation such as conspiracies, discrediting, emotion, or social media feeds that may influence democratic processes [6]. In addition, manual contents on the web are intentionally polluted by fake journalists profiling to attract clicks and attention [7]. In addition, it is very important to evaluate the credibility of the writers' beliefs and moral values.

This paper focuses on analyzing the performance of several models for fake news detection in Romanian by using neural network architectures such as long short-term memory (LSTM), a convolutional neural network (CNN), gated recurrent units (GRU), Bidirectional Encoder Representations from Transformers (BERT), and standard classifiers such as Support Vector Machine (SVM) and Naïve Bayes (NB). Furthermore, in this paper, a statistical analysis for a dataset of online articles with real and fake classes is presented. Moreover, a sentiment analysis and irony detection systems are applied to our datasets, providing important information in the process of detecting fake news.

The paper continues as follows. In the next section, the related work is presented. Section 3 describes the methodology of the proposed method, and Section 4 describes the experiments, while Section 5 presents the results. Section 6 contains the conclusions.

2. Related Work

Busioc et al. [8] proposed an automated analysis of political statements in Romanian using several natural language processing techniques. They used a corpus collected from Factual [9], a Romanian initiative where different public statements which are labeled as true or fake can be found. Another study highlights the existing approaches, challenges, and observations from other languages to be applied for Romanian resources, identifying future paths [10] in developing fake news detection systems.

In recent years, there have been many studies available in the fake news detection field for other languages. Currently, there are studies and models that suggest using classical machine learning algorithms for detecting fake news in other languages [11]. Some authors have also proposed BERT models [12], mentioning an accuracy of 98.90% for FakeBert. There is a wide choice of deep neural network models available in the literature [13], and some papers use hybrid convolutional neural network and recurrent neural network (RNN) models, such as that of Ajao et al. [14], wherein they achieved an accuracy of 82%. Furthermore, several papers focused on fake news detection using neural learning systems have been published [15–17], highlighting the complexity of this domain. Another study provides a comparison between multiple methods using neural network systems and attention mechanisms [18], achieving an accuracy of 88.78% for CNN + Bi-LSTM ensembled networks.

The internet gives researchers the opportunity to find and use several datasets, such as Fake News Challenge [19], to develop different approaches based on CNN, LSTM, and Bi-LSTM [20] by using the headlines and bodies of the articles, achieving an accuracy of 71.2% for the testing dataset.

At this moment, the importance of social media is widely known from the social to the financial points of view, being a powerful free online tool for spreading misinformation without investigation or personal filters. Due to the increasing number of people who collected from social media their daily information, these platforms became the most important “weapon” in misinformation campaigns. Recent papers [21] have revealed that fake and real news are spreading differently and deeply, making it possible to create patterns that could be used in fake news detection systems.

An existing paper in the broader literature examined the challenges in fake news detection on social media using a logistic classifier [22], achieving an accuracy of 99.4% for the testing data.

The study of Shu et al. [23] revealed that the user's engagements are important and relevant. In addition, the analysis process of the huge amount of noisy data collected from social platforms is essential in understanding and detecting the misinformation phenomenon.

Guibon et al. [24] proposed different approaches for fake news detection systems and, based on redundant information, tried to find a connection between satire and fake news, achieving an accuracy of 93% for several datasets.

Some experiments that have been used to incorporate sentiment analysis (SA) in fake news detection approaches were presented by Zhang et al. [25], who stated that most existing papers on fake news are based on the strength of the emotions expressed by the publishers. Additionally, Ross and Thirunarayan [26] used in their study several sentimental features.

Nowadays, large-scale pre-trained language models have become very important in fake news detection, and the first Romanian transformer-based language model was proposed by Dumitrescu, Avram, and Pyysalo [27]. In this paper, a Romanian pretrained BERT model for experiments named RoBERT is used [28].

3. Methodology

3.1. Dataset Details

We collected a dataset of fake and real news between 2016 and 2021:

- Fake news: This dataset contains 12,767 news items, and it was automatically crawled from Romanian online platforms such as Fluierul [29], Vremuritulburi [30], and Cunoastelumea [31]. It is based on Rubrika [32], the first fully automatic news aggregator in Romania, which promotes articles only from trustworthy sources and provides a list of websites to avoid [33]. In addition to this automatically collected dataset, there were 297 more news items added that were manually labeled as fake news. For example, after a fake news instance was manually annotated, in the Romanian online environment, several news sites with the same information that was already propagated can be identified, and these news sites are added to the dataset by a human using a web application system, labeling them with a simple button as fake.
- Real news: This dataset contains 25,841 news items and was manually collected from Romanian official sources such as Agerpres [34], Mediafax [35], and Rador [36]. Each article's content is verified and annotated as real news.

The dataset includes only Romanian content that was automatically detected with a PHP library named Text Language Detect [37]. The evaluation and annotation process was performed by 12 employees, consisting of males and females aged between 34 and 49 years, in a public institution in Romania.

3.2. Dataset Description

For the research presented herein, the dataset was split into train (59.48%), validation (20.26%), and test (20.26%). Additionally, to maintain the balance of class distribution, only 50.55% of the real class was used for this experiment, being randomly selected. Table 1 presents the distribution of the fake and real classes in the dataset. It shows that our dataset was balanced, and the training dataset was three times longer than the validation and test datasets.

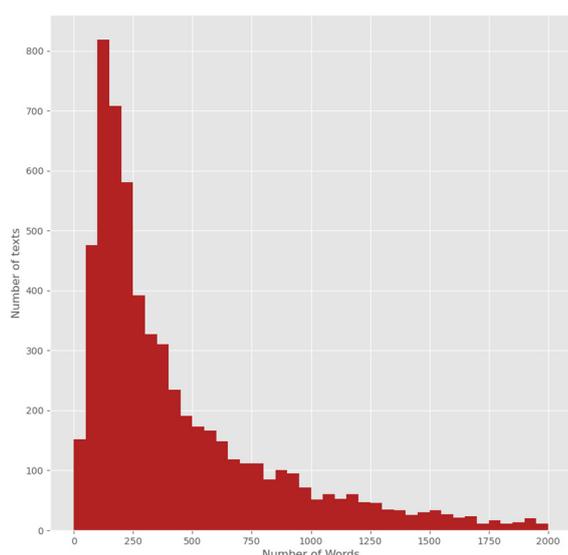
Table 1. Distribution of datasets across classes.

Dataset	Fake	Real	Total
Training	7768	7768	15,536
Validation	2648	2648	5296
Test	2648	2648	5296
Total	13,064	13,064	26,128

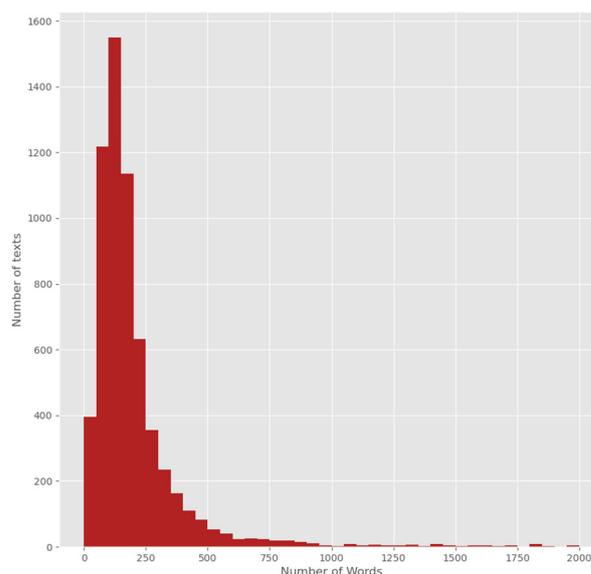
Table 2 presents two examples of real and fake news, and Figure 1 presents the distribution of words across the fake and real datasets. Table 3 shows that the average number of words of the fake dataset was three times larger than that of the true dataset, and the vocabulary size consisting of unique Romanian words from The Explanatory Dictionary of the Romanian Language (DEX) for the fake dataset was twice as large as that of the real dataset.

Table 2. Examples of fake and real news from dataset.

Label	Source	Text
Fake	constantadeazi.ro	<p>Militar român rănit în Afganistan, trecut în rezervă cu o pensie de 360 lei. “Am stat trei luni în comă, am suferit 85 de fracturi. Dar atât costă viața unui militar” . . .</p> <p>(English translation: Romanian soldier injured in Afghanistan, retired with a 360 lei pension. “I was in a coma for three months and suffered 85 fractures. But this is how much a soldier’s life is worth” . . .)</p> <p>https://www.constantadeazi.ro/militar-roman-ranit-in-afghanistan-trecut-in-rezerva-cu-o-pensie-de-360-lei-am-stat-trei-luni-in-coma-am-suferit-85-de-fracturi-dar-atat-costa-viata-unui-militar</p>
True	adevarul.ro	<p>MAPN a demontat „povestea tragică” a eroului impostor de la „Chefi la cuțite”, rănit în Afganistan și trecut în rezervă cu o pensie mica . . .</p> <p>(English translation: MoND dismantled the “tragic story” of the impostor hero from “Chefi la cuțite” who was injured in Afghanistan and retired with a small pension . . .)</p> <p>https://adevarul.ro/entertainment/tv/mapn-demontat-povestea-tragica-eroului-impostor-chefi-cutite-ranit-afagnistan-trecut-rezerva-pensie-mica-spune-antena-1-motivul-nu-l-elimina-1_5aeb1b3edf52022f758a8c01/index.html</p>



(a)



(b)

Figure 1. Distribution of words: (a) fake dataset and (b) real dataset.

Table 3. Distribution of words across dataset.

Words	Fake	Real
Romanian unique words	102,006	44,969
Average words per news	413	155

In some cases, the differences between fake and real news could be rather small, and the presented example was very hard to identify because it was necessary to search the employee in the public institution archive when he said that he was injured as a soldier, which turned out to be a lie.

3.3. Data Pre-Processing

Before applying machine learning algorithms and neural network architectures, data pre-processing is a necessary task. The Romanian stop words (e.g., “acea”, “apoi”, “atare”, etc.) and punctuation marks (e.g., “!”, “?”, “-”, etc.) were removed and HTML tags were eliminated using Python libraries such as Pandas and NumPy.

Pre-processing was applied to minimize noisy data and provide simple, complete, and consistent datasets. In order, the words of the fake and real training datasets with more than 5000 occurrences were removed. Some examples of such words included “Romania” (15,461), “arta” (14,769), “national” (13,729), “military” (19,111), and “present” (8085).

This paper proposes a method that receives as input pre-processed news to reduce the chances of underfitting or overfitting and performs several analyses and transformations, using the term frequency-inverse document frequency (TF-IDF) for feature extraction.

4. Experiments

For the experiments presented in this paper, 4 NVIDIA Tesla V100 GPU Accelerators with 32 GB RAM and 5120 CUDA cores were used. The proposed method (as shown in Figure 2) consisted of two classical algorithms (NB and SVM), three deep learning models (LSTM, CNN, and GRU), and two variants of BERT.

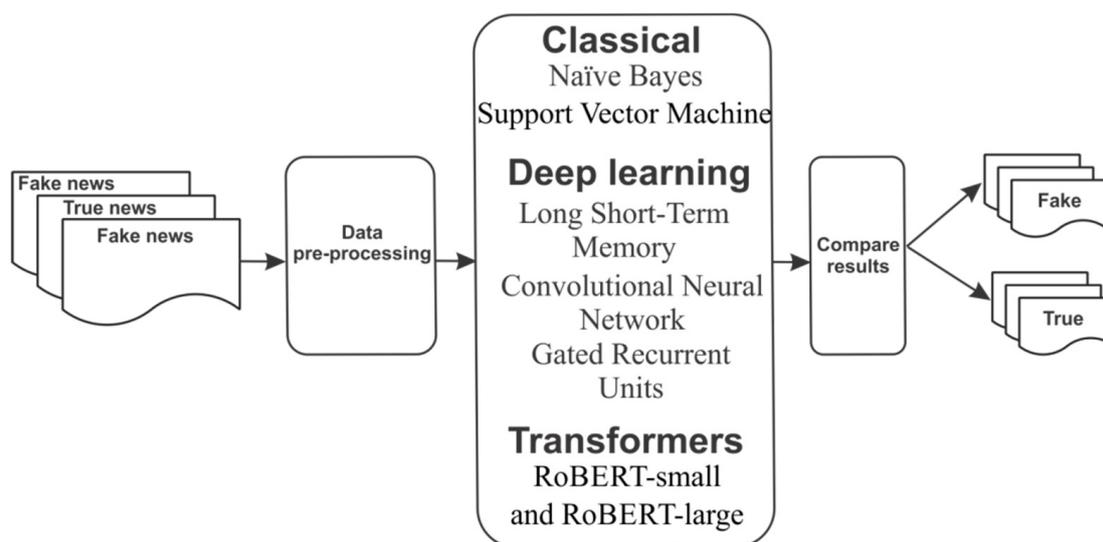


Figure 2. Architecture of the proposed fake news detection system.

4.1. Proposed Models

- **Classical algorithms:** The classical machine learning models are based on supervised classifiers such as Naïve Bayes and Support Vector Machine. Each traditional algorithm learns in different ways. The Naïve Bayes algorithm is based on Bayes’ theorem to evaluate and choose the highest probability of new data belonging to one of the classes

defined in the dataset. The SVM classifier finds the best hyperplane that separates the data into two classes (fake vs. real) with the highest margin. For experiments, the SVM algorithm uses an SVC linear kernel, and the NB algorithm uses multinomial Naive Bayes.

- Deep learning models: Three types of deep neural network models were investigated. The first two were recurrent neural network architectures using LSTM and GRU. The third type was a CNN architecture that is a class of deep neural network mostly used in computer vision tasks. For the experiments, these architectures used the optimal parameters achieved during the random search optimization phase and binary cross-entropy as a loss function.
- Transformer models: Transformers are a type of neural network model, being introduced by Vaswani et al. [38] to solve the issue of sequence transduction or neural machine translation. The most popular NLP model that uses a transformer is BERT, introduced by Devlin et al. [39], which is a model that learns contextual embeddings from both sides of a token's context during the training phase.

The two applications of BERT are “pretraining” and “fine-tuning”. For the pretraining process, BERT uses the masked language model and next sentence prediction. In this research, for the fine-tuning process, the Romanian pretrained model RoBERT was used. Currently, there are three uncased versions available: RoBERT-small, RoBERT-base, and RoBERT-large.

4.2. Deep Learning Architectures

This section presents several pieces of software or packages and version numbers, as shown in Table 4, and it also provides an overview of the deep neural network models used in this research, the main hyperparameters (as shown in Table 5), and their architectures, being fine-tuned for three epochs with an Adam optimizer.

Table 4. Software or packages and version numbers.

Software or Packages	Version
CUDA	11.2
Python	3.8.5
Keras	2.4.3
Nltk	3.5
TensorFlow-GPU	1.14.0

Table 5. Hyperparameters of deep learning models.

Parameter Name	Value of Parameter
Learning rate	0.001
Neurons	32
Optimizer	Adam
Weights	random
Dropout	0.2
Batch size	32
Vocabulary size	50,000
Number of words	5000

- Long short-term memory: LSTM networks are a type of recurrent neural network having the capability to learn a mapping between the input and output patterns. For the experiments, the LSTM model consisted of 1 layer with 128 units that decreased the embedding vector from 5000 to 128, a dropout layer (0.2), and 2 dense layers, using 32 as the batch size and 32 neurons. The details of the LSTM architecture used in this work are presented in Table 6.

Table 6. LSTM layered architecture.

Label	Output Size	Param Number
Embedding	5000×32	1,600,000
LSTM	128	82,432
Dropout	128	0
Dense	5	645
Dense	1	6

- Convolutional neural network: A CNN is a deep learning architecture successfully used to extract features for images and classify text documents. For this architecture, the convolution layer has 250 filters with a kernel size of 3 that decreases the embedding vector from 5000 to 4998. A max-pooling, Rectified Unit Layer (RELU), activation, and dropout layer were added to the proposed CNN model, passing the outputs through a dense layer. The CNN architecture is described in Table 7.

Table 7. CNN layered architecture.

Label	Output Size	Param Number
Embedding	5000×32	1,600,000
Conv1D	4998×250	24,250
Maxpool1D	250	0
Activation	250	0
Dropout	250	0
Dense	1	251

- Gated recurrent units: GRU are one of the latest generation of recurrent neural networks, being more complex due to a hidden state which transfers useful information based on two gates: a reset gate and an update gate. In this architecture, the GRU model consists of one layer with 128 units and a dropout, activation (TanH, the hyperbolic tangent), and dense layer. The detail of the GRU architecture that is used in this work is presented in Table 8.

Table 8. GRU layered architecture.

Label	Output Size	Param Number
Embedding	5000×32	1,600,000
GRU	128	62,208
Activation	128	0
Dropout	128	0
Dense	1	129

4.3. Transformer Architectures

The transformer model is an encoder-decoder architecture using a multi-headed attention layer to increase the speed of the training process and excel in specific NLP tasks such as voice conversion or text-to-speech transformation.

In this research, two versions of Romanian pretrained BERT models were used: RoBERT-small and RoBERT-large. RoBERT is a Romanian pretrained model that is based on a multi-layer bidirectional transformer. It consists of a large Romanian corpus collected from several sources such as Wikipedia, Oscar [40], and the RoTex collection [41]:

- RoBERT-small (see Table 9) contains less weights (19M) and a number of trainable layers.
- RoBERT-large (see Table 9) contains large weights (341M) and twice the number of trainable layers, having the same layer sizes as BERT-large.

Table 9. RoBERT models: weights in millions (W), vocabulary size (V), number of layers (L), hidden size (H), and number of attention heads (A).

Model	W	V	L	H	A
RoBERT-small	19M	38,000	12	256	8
RoBERT-large	341M	38,000	24	1024	16

Furthermore, the authors of the RoBERT models followed the same methodology proposed by Devlin et al. to train their models.

The BERT models were trained over 30 epochs with the Adam optimizer [42], having a learning rate of 3×10^{-5} and maximum sequence length of 512. Moreover, these models contained two dense layers: a dropout layer (0.2) with RELU activation, and a Softmax layer. Table 10 presents the main hyperparameters used in the training process.

Table 10. Hyperparameters of BERT models.

Parameters Name	Value of Parameter
Number of epochs	30
Batch size	32
Optimizer	Adam
Loss function	Categorical cross-entropy
Dropout	0.2
Learning rate	0.00003

5. Results and Discussion

This paper presented three architectures based on classical algorithms, deep learning models, and transformers. Tables 11 and 12 show the results of the test and validation datasets, which consisted of 5296 unique news items.

Table 11. Scores for classical, deep learning, and transformer models for validation dataset (5296 news items).

Model Type	Model Name	Acc	Pre	Rec	F1	TP	TN	FP	FN
Classical	SVM	0.944	0.904	0.988	0.944	2519	2480	266	31
	NB	0.976	0.957	0.994	0.975	2534	2633	113	16
Deep learning	LSTM	0.967	0.939	0.997	0.967	2542	2581	165	8
	CNN	0.978	0.965	0.991	0.978	2528	2654	92	22
	GRU	0.961	0.927	0.997	0.961	2543	2545	201	7
Transformers	RoBERT-small	0.933	0.896	0.975	0.934	2485	2457	289	65
	RoBERT-large	0.907	0.851	0.976	0.910	2490	2311	435	60

Table 12. Scores for classical, deep learning, and transformer models for test dataset (5296 news items).

Model Type	Model Name	Acc	Pre	Rec	F1	TP	TN	FP	FN
Classical	SVM	0.945	0.912	0.985	0.947	2609	2397	251	39
	NB	0.975	0.964	0.986	0.975	2610	2551	97	38
Deep learning	LSTM	0.979	0.964	0.994	0.979	2633	2551	97	15
	CNN	0.981	0.971	0.992	0.982	2628	2569	79	20
	GRU	0.975	0.958	0.992	0.975	2627	2534	114	21
Transformers	RoBERT-small	0.919	0.866	0.992	0.925	2628	2240	408	20
	RoBERT-large	0.870	0.816	0.955	0.880	2529	2078	570	119

- **Classical algorithms:** From Table 12, it can be observed that the Naïve Bayes algorithm obtained a better F1 score of 97.50% for the test dataset compared with the Support Vector Machine algorithm, which obtained an F1 score of 94.70%. The results were slightly similar for the validation set. There are studies and models that suggest using Naïve Bayes with n-gram (bigram TF-IDF) features to outperform the standard machine learning systems for online fake news detection approaches, achieving almost 94% accuracy on multiple corpora [43].
- **Deep learning models:** In this research, the differences between the neural network models' performances were small. The CNN architecture obtained an F1 score of 97.80% for the validation dataset and an F1 score of 98.20% for the test dataset, outperforming the LSTM and GRU models. For example, instead of just using CNN models, another study [44] proposed a hybrid deep learning architecture that combines the CNN and RNN models trained on several datasets.
- **Transformer models:** As already mentioned, this research used for the BERT experiments two Romanian pretrained models. The RoBERT-small model obtained a better F1 score of 92.50%, while RoBERT-large's was only 88% for the test dataset, achieving similar results for the validation dataset. This was due to the first dense layer of the BERT models, which decreased the dense vector from 1024 to 512 for RoBERT-large and from 256 to 32 for RoBERT-small, indicating that the RoBERT-small model was more efficient for our datasets, generating fewer false positives. Future research should consider the potential effects of Language Understanding with Knowledge-Based Embeddings (LUKE), a new model based on the transformer that outperformed the BERT and RoBERTa [45] models, achieving an F1 score of 95%. LUKE [46] is based on the Stanford Question Answering Dataset [47].
- **Fake news and sentiment analysis:** The sentiment expressed in the fake news dataset had a significant role, and some researchers such as Alonso et al. [48] and Bhutani et al. [49] proposed different fake news detection systems that incorporated sentiment as an important feature. Therefore, a sentiment analysis method [50] was applied to the test dataset, based on an algorithm that achieved an F1 score of 82% using a Romanian dictionary of 42,497 labeled words with 3 levels for the positive and negative polarities. Table 13 shows that 99.96% of the fake news dataset contained a neutral polarity, indicating that in these campaigns of fake news, the impartial connotation was predominant. Moreover, Figure 3 presents as a word cloud several words with positive (left side) and negative polarities (right side) from the fake news employed in the proposed system.

important. This paper supports these statements in defining the necessary environment that would lead to the identification and processing of available data, which will help us to control this phenomenon.

In this paper, a fake news detection method was proposed by choosing the best result between traditional classifiers, such as the Naïve Bayes, Support Vector Machine, and neural network models, based on long short-term memory, convolutional neural networks, gated recurrent units, and BERT models. The proposed approach achieved a better score for the convolutional neural network, indicating the superiority of deep neural network architectures over machine learning. Furthermore, according to the results that were obtained during the evaluation process, by applying an SA and irony detection system, a correlation could be associated between irony or neutral sentiments and fake news in Romanian online news. In addition, our experiments confirmed that human evaluation provides paramount contributions for Romanian datasets and achieves significant improvements for fake news detection approaches when using a convolutional neural network. Our results for the Romanian language are consistent with the literature covering English news, which shows that some studies investigated and analyzed several datasets using neural network architectures, achieving better performance for the convolutional neural network (FNDNet) [53] and yielding an accuracy of 98.36% for the test data. Moreover, the paper of Martínez-Gallego [54] addressed the problem of fake news detection, achieving an accuracy of up to 80% for a language with Latin roots such as Spanish by using a combination of a pretrained BERTO (Spanish BERT) model and an LSTM architecture.

The system proposed in this paper was integrated in a mobile application [55] designed and developed to give users access to reliable Romanian information and to inspire other researchers to use this model.

Nowadays, governments and public institutions are using different machine learning algorithms to automate the important processes of several departments, such as human resources and claims revenue, with it being necessary to follow several ethical rules such as transparency, responsibility, or fairness and to avoid epistemic and normative concerns. The exclusive use of algorithms and AI systems may encounter errors or have insufficient data for selecting the best ways to work, but these systems should deliver socially good outcomes, and ethical and technological analyses are mandatory.

Author Contributions: Conceptualization, M.C.B.; methodology, M.C.B.; software, M.C.B.; validation, S.T.-M. and T.R.; investigation, M.C.B., S.T.-M., and T.R.; resources, M.C.B., S.T.-M., and T.R.; writing—original draft, M.C.B.; writing—review and editing, S.T.-M. and T.R.; visualization, M.C.B. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Acknowledgments: We would like to thank the authors of all datasets used in this paper for making the data available to the community.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Brown, T.B.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. Language Models are Few-Shot Learners. In Proceedings of the 33th Conference on Neural Information Processing Systems (NeurIPS 2020), Vancouver, BC, Canada, 6–12 December 2020; pp. 1877–1901.
2. Althuis, J.; Haiden, L. *Fake News: A Roadmap*; NATO Strategic Communications Centre of Excellence: Riga, Latvia, 2018.
3. García, A.S.; García, G.G.; Prieto, S.M.; Guerrero, A.J.M.; Jiménez, R.C. The Impact of Term Fake News on the Scientific Community. *Sci. Perform. Mapp. Web Sci. Soc. Sci.* **2020**, *9*, 73. [[CrossRef](#)]
4. Nordberg, P.; Kävrestad, J.; Nohlberg, M. Automatic detection of fake news. In Proceedings of the 6th International Workshop on Socio-Technical Perspective in IS Development, Grenoble, France, 8–9 June 2020; pp. 168–179.

5. Caplan, R.; Hanson, L.; Donovan, J. Dead reckoning: Navigating Content Moderation after “Fake News”. Data & Society Research Institute. 2018. Available online: <https://datasociety.net/output/dead-reckoning> (accessed on 26 September 2021).
6. Van der Linden, S. Beating the hell out of fake news. *Ethical Rec. Proc. Conway Hall Ethical Soc.* **2017**, *122*, 4–7.
7. Meel, P.; Vishwakarma, D.K. Fake news, rumor, information pollution in social media and web: A contemporary survey of state-of-the-arts, challenges and opportunities. *Expert Syst. Appl.* **2020**, *153*, 112986. [CrossRef]
8. Busioc, C.; Dumitru, V.; Ruseti, S.; Terian-Dan, S.; Dascalu, M.; Rebedea, T. What Are the Latest Fake News in Romanian Politics? An Automated Analysis Based on BERT Language Models. In *Ludic, Co-design and Tools Supporting Smart Learning Ecosystems and Smart Education: Proceedings of the 6th International Conference on Smart Learning Ecosystems and Regional Development, Bucharest, Romania, 24–25 June 2021*; Springer Nature: Singapore, 2021; Volume 249, p. 201. [CrossRef]
9. Factual. Available online: <https://www.factual.ro> (accessed on 26 September 2021).
10. Busioc, C.; Ruseti, S.; Dascalu, M. A Literature Review of NLP Approaches to Fake News Detection and Their Applicability to Romanian-Language News Analysis. *Transilv. J.* **2020**, 65–71. [CrossRef]
11. Manzoor, S.I.; Nikita, J.S. Fake News Detection Using Machine Learning approaches: A systematic Review. In Proceedings of the 3rd International Conference on Trends in Electronics and Informatics (ICOEI), Tirunelveli, India, 23–25 April 2019; pp. 230–234.
12. Kaliyar, R.K.; Goswami, A.; Narang, P. FakeBERT: Fake news detection in social media with a BERT-based deep learning approach. *Multimed. Tools Appl.* **2021**, *80*, 11765–11788. [CrossRef] [PubMed]
13. Wang, W.Y. Liar, Liar Pants on Fire: A New Benchmark Dataset for Fake News Detection. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, Vancouver, BC, Canada, 30 July–4 August 2017; Volume 2, pp. 422–426. [CrossRef]
14. Ajao, O.; Deepayan, B.; Shahrzad, Z. Fake News Identification on Twitter with Hybrid CNN and RNN Models. In Proceedings of the 9th International Conference on Social Media and Society, Copenhagen, Denmark, 18–20 July 2018; pp. 226–230. [CrossRef]
15. Jiang, T.; Li, J.P.; Haq, A.U.; Saboor, A.; Ali, A. A Novel Stacking Approach for Accurate Detection of Fake News. *IEEE Access* **2021**, *9*, 22626–22639. [CrossRef]
16. Saikh, T.; De, A.; Ekbal, A.; Bhattacharyya, P. A deep learning approach for automatic detection of fake news. In Proceedings of the 16th International Conference on Natural Language Processing, Hyderabad, India, 18–21 December 2019; pp. 230–238.
17. Thota, A.; Tilak, P.; Ahluwalia, S.; Lohia, N. Fake News Detection: A Deep Learning Approach. *SMU Data Sci. Rev.* **2018**, *1*, 10. Available online: <https://scholar.smu.edu/datasciencereview/vol1/iss3/10> (accessed on 26 September 2021).
18. Kumar, S.; Asthana, R.; Upadhyay, S.; Upreti, N.; Akbar, M. Fake news detection using deep learning models: A novel approach. *Trans. Emerg. Telecommun. Technol.* **2020**, *31*, e3767. [CrossRef]
19. Fake News Challenge. Available online: <http://www.fakenewschallenge.org> (accessed on 26 September 2021).
20. Abedalla, A.; Al-Sadi, A.; Abdullah, M. A Closer Look at Fake News Detection: A Deep Learning Perspective. In Proceedings of the 3rd International Conference on Advances in Artificial Intelligence, Istanbul, Turkey, 26–28 October 2019; pp. 24–28. [CrossRef]
21. Talwar, S.; Dhir, A.; Singh, D.; Virk, G.S.; Salo, J. A Sharing of fake news on social media: Application of the honeycomb framework and the third-person effect hypothesis. *J. Retail. Consum. Serv.* **2020**, *57*, 102197. [CrossRef]
22. Aldwairi, M.; Alwahedi, A. Detecting Fake News in Social Media Networks. In Proceedings of the 9th International Conference on Emerging Ubiquitous Systems and Pervasive Networks, Leuven, Belgium, 5–8 November 2018; Volume 141, pp. 215–222. [CrossRef]
23. Shu, K.; Sliva, A.; Wang, S.; Tang, J.; Liu, H. Fake News Detection on Social Media: A Data Mining Perspective. *ACM SIGKDD Explor. Newsl.* **2017**, *19*, 22–36. [CrossRef]
24. Guibon, G.; Ermakova, L.; Seffih, H.; Firsov, A.; Noé-Bienvenu, G.L. Multilingual Fake News Detection with Satire. In Proceedings of the International Conference on Computational Linguistics and Intelligent Text Processing, La Rochelle, France, 7–13 April 2019.
25. Zhang, X.; Cao, J.; Li, X.; Sheng, Q.; Zhong, L.; Shu, K. Mining Dual Emotion for Fake News Detection. In Proceedings of the Web Conference 2021, World Wide Web Conference, Ljubljana, Slovenia, 19–23 April 2021; ACM: New York, NY, USA, 2021; pp. 3465–3476. [CrossRef]
26. Ross, J.; Thirunarayan, K. Features for Ranking Tweets Based on Credibility and Newsworthiness. In Proceedings of the 2016 International Conference on Collaboration Technologies and Systems, Orlando, FL, USA, 31 October–4 November 2016; pp. 18–25. [CrossRef]
27. Dumitrescu, S.D.; Avram, A.; Pyysalo, S. The birth of Romanian BERT. In *Findings of the Association for Computational Linguistics, Proceedings of the EMNLP 2020, Online Conference, 16–20 November 2020*; Association for Computational Linguistics: Stroudsburg, PA, USA, 2020; pp. 4324–4328. [CrossRef]
28. Masala, M.; Ruseti, S.; Dascalu, M. RoBERT—A Romanian BERT Model. In Proceedings of the 28th International Conference on Computational Linguistics, International Committee on Computational Linguistics, Barcelona, Spain, 8–13 December 2020; pp. 6626–6637. [CrossRef]
29. Fluierul. Available online: <https://www.fluierul.ro> (accessed on 26 September 2021).
30. Vremuritulburi. Available online: <https://vremuritulburi.com> (accessed on 26 September 2021).
31. Cunoastelumea. Available online: <https://www.cunoastelumea.ro> (accessed on 26 September 2021).
32. Rubrika. Available online: <https://rubrika.ro> (accessed on 26 September 2021).

33. Mediafax & Rubrika. Available online: <https://www.mediafax.ro/life-inedit/aplicatia-de-necrezut-vaneaza-fake-news-urile-din-romania-care-sunt-cele-70-de-site-uri-vizate-16156645> (accessed on 26 September 2021).
34. Agerpres. Available online: <https://www.agerpres.ro> (accessed on 26 September 2021).
35. Mediafax. Available online: <https://mediafax.ro> (accessed on 26 September 2021).
36. Rador. Available online: <https://rador.ro> (accessed on 26 September 2021).
37. Text Language Detect. Available online: https://pear.php.net/package/Text_LanguageDetect (accessed on 26 September 2021).
38. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, L.; Polosukhin, I. Attention is all you need. In Proceedings of the Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems, Long Beach, CA, USA, 4–9 December 2017; pp. 5998–6008.
39. Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT, Minneapolis, MN, USA, 2–7 June 2019; Volume 1, pp. 4171–4186. [CrossRef]
40. Suarez, P.O.; Sagot, B.; Romary, L. Asynchronous Pipeline for Processing Huge Corpora on Medium to Low Resource Infrastructures. In Proceedings of the 7th Workshop on the Challenges in the Management of Large Corpora (CMLC-7), Cardiff, UK, 22 July 2019; Leibniz-Institut für Deutsche Sprache: Mannheim, Germany, 2019; pp. 9–16. [CrossRef]
41. RoTex Corpus Builder-Builds a Corpus of Romanian Text, Suitable for NLP Research, from Different Online Sources. Available online: <https://github.com/aleris/ReadME-RoTex-Corpus-Builder> (accessed on 26 September 2021).
42. Kingma, D.; Ba, J. Adam: A Method for Stochastic Optimization. International Conference on Learning Representations. *arXiv* **2015**, arXiv:1412.6980.
43. Khan, J.Y.; Khondaker, T.I.; Afroz, S.; Uddin, G.; Iqbal, A. A benchmark study of machine learning models for online fake news detection. *Mach. Learn. Appl.* **2021**, *4*, 100032. [CrossRef]
44. Nasir, J.A.; Khan, O.S.; Varlamis, I. Fake news detection: A hybrid CNN-RNN based deep learning approach. *Int. J. Inf. Manag. Data Insights* **2021**, *1*, 100007. [CrossRef]
45. Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; Stoyanov, V. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv* **2019**, arXiv:1907.11692.
46. Yamada, I.; Asai, A.; Shindo, H.; Takeda, H.; Matsumoto, Y. LUKE: Deep Contextualized Entity Representations with Entity-aware Self-attention. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, Online Conference, 16–20 November 2020; pp. 6442–6454. [CrossRef]
47. Rajpurkar, P.; Zhang, J.; Lopyrev, K.; Liang, P. SQuAD: 100,000+ Questions for Machine Comprehension of Text. In Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, Austin, TX, USA, 1–5 November 2016; pp. 2383–2392. [CrossRef]
48. Alonso, M.A.; Vilares, D.; Gómez-Rodríguez, C.; Vilares, J. Sentiment Analysis for Fake News Detection. *Electronics* **2021**, *10*, 1348. [CrossRef]
49. Bhutani, B.; Rastogi, N.; Sehgal, P.; Purwar, A. Fake News Detection Using Sentiment Analysis. In Proceedings of the Twelfth International Conference on Contemporary Computing (IC3), Noida, India, 8–10 August 2019; pp. 1–5. [CrossRef]
50. Buzea, M.C.; Trăușan-Matu, Ș.; Rebedea, T. A Three Word-Level Approach Used in Machine Learning for Romanian Sentiment Analysis. In Proceedings of the 18th RoEduNet Conference: Networking in Education and Research, Galati, Romania, 10–12 October 2019; pp. 1–6. [CrossRef]
51. Buzea, M.C.; Trăușan-Matu, Ș.; Rebedea, T. Automatic Irony detection for Romanian online news. In Proceedings of the 24th International Conference on System Theory, Control and Computing (ICSTCC), Sinaia, Romania, 8–10 October 2020; pp. 72–77. [CrossRef]
52. Times New Roman. Available online: <https://www.timesnewroman.ro> (accessed on 26 September 2021).
53. Kaliyar, R.K.; Goswami, A.; Narang, P.; Sinha, S. FNDNet—A deep convolutional neural network for fake news detection. *Cogn. Syst. Res.* **2020**, *61*, 32–44. [CrossRef]
54. Martínez-Gallego, K.; Álvarez-Ortiz, A.M.; Arias-Londoño, J.D. Fake news detection in spanish using deep learning techniques. *arXiv* **2021**, arXiv:2110.06461.
55. Buzea, M.C.; Trăușan-Matu, Ș.; Rebedea, T. Targeted Romanian Online News in a Mobile Application Using AI. In Proceedings of the RoCHI-International Conference on Human-Computer Interaction, Sibiu, Romania, 22–23 October 2020; pp. 54–60. [CrossRef]