Romanian Fake News Identification using Language Models

Andrei Preda
University Politehnica of Bucharest
313 Splaiul Independetei, Bucharest, Romania
andrei.preda3006@stud.acs.upb.ro

Stefan Ruseti
University Politehnica of Bucharest
313 Splaiul Independetei, Bucharest, Romania
stefan.ruseti@upb.ro

Simina-Maria Terian
Lucian Blaga University of Sibiu
5-7 Victoriei Blvd.
Sibiu, Romania
simina.terian@ulbsibiu.ro

Mihai Dascalu
University Politehnica of Bucharest
313 Splaiul Independetei, Bucharest, Romania
mihai.dascalu@upb.ro

ABSTRACT
In an increasingly complex socio-economic and political context, the amount of fake news distributed online is on the rise and has already influenced major events and our decision-making capabilities. Studies show that people tend to be overconfident in their ability to identify fake news, which suggests that an automatic system for detecting them might be helpful. This article describes state-of-the-art techniques used in text classification and analyzes the performance of different neural networks on a corpus of news articles written in Romanian. Classical machine learning methods are considered, as well as more complex models based on Transformers, which achieved better results, having a weighted F1-score of .75 using RoBERT and CNN on top. Experiments with multi-task learning are also described but did not provide a boost in performance while reaching an F1-score of .74. We also introduce a prototype web application and additional use cases for automated fake news detection systems.

Author Keywords
Fake News; Natural Language Processing; Text Classification; BERT; Multi-task Learning; Romanian language.

ACM Classification Keywords
I.2.7 Natural Language Processing: Text analysis.

DOI: 10.37789/rochi.2022.1.1.*

INTRODUCTION
In a political context that seems to become more unpredictable by the day, the fake news phenomenon appears to be gaining momentum and affecting more and more parts of the world. With authoritarian governments restricting the information their population is exposed to, access to honest, unbiased news is very precious.

While the concept of fake news is familiar to most people at an intuitive level, it can be quite hard to precisely describe it. We could say that fake news is fabricated information that mimics normal news [11], but which also aims either to misinform, cause harm, or push a political agenda. In any case, the producers and spreaders of fake news “attempt to determine the recipient to perform a certain type of action” [16]. As such, fake news can come in many forms, either as fabricated articles, propaganda pieces, or information taken out of context, maybe even presented with a fake context. Even parodies or satire could be read in certain contexts as false news, because, even though they do not intend to harm, they could still "fool the reader" [20].

In recent years, the proliferation of fake news has brought more attention to the phenomenon from the public. The Internet has allowed fake news to flourish, for example by sensationalistic content being shared by users on social media. In fact, the very act of sharing represents a new kind of speech act, without any clear equivalent in the traditional media and, for this reason, with fuzzy moral standards and borders [13]. Given this communicational and ethical ambiguity, entire websites dedicated to producing and/or spreading manipulative content have appeared. As a result, fact-checking websites have appeared too, such as Veridica[1].

Fake news which gains traction and attracts many readers have the potential to produce important changes in society. For example, numerous studies have shown that political elections over the world have been swayed by fake news, as proven by the cases of the United States [4], Portugal [3], Brazil [2], Nigeria [12], and Taiwan [19].

Concerningly, some studies show that people tend to overestimate their ability to recognize fake news [8, 15]. This is confirmed, on the one hand, by the fact that they often label as fake news any opinion that challenges their partisan beliefs [18] and, on the other, by the obvious affinity between fake news and conspiracy theories [9]. Both phenomena indicate the increasing difficulty of reaching unbiased perspectives in

[1] https://www.veridica.ro/

the contemporary world and suggest that automatic systems for checking the legitimacy of news could be beneficial and are worth trying to develop.

Such automatic systems could help users avoid fake news in multiple ways. For example, user-facing tools could be developed to either warn about articles that seem suspicious or outright block them. On the other hand, they could also serve as internal tools for websites that practice news curation. These systems could also help journalists who want to assess the credibility of a news source, or other people who have to deal with a large amount of information whose authenticity is not guaranteed.

METHOD

Corpus

This study considers the Fakerom dataset[2] that contains around 14,000 articles centered mainly around the subject of COVID-19. Out of these, approximately 1,200 are labeled with article types and can be used for fake news detection. The annotation process was done by experts.

Each article was classified in one of six categories: *real* news, *plausible* (or *authentic*) news, *propaganda* pieces, *fabricated* news, *satire* and *fictional*. The six categories can be grouped into pairs and thus they cover the areas of *true* (= real + plausible/authentic), *fake* (= propaganda + fabricated), and *imaginary* news (= satirical + fictional) [17].

Although possible, it is not probable that a source publishes fake news all the time, since it could soon lose its credibility and its readership. As such, we assume that sources publish a combination of different types of articles, with some of the categories being predominant. Our articles were scraped from approximately 87 websites, and 38 of these sources provided multiple types of articles.

Apart from being labeled, the articles were also annotated with other metadata, such as general categories in which the subject of the article belongs (such as health, politics. or religion), the most important action described in the article, or keywords and comments.

Some articles required data cleaning because they were scraped from news websites and still contained artifacts from these websites. For example, a batch of articles started with a fixed prefix that appeared to be part of the web page. Others contained sentences that represented metadata such as the author, the time when the article was published, sources (mostly representing URLs), or sentences meant to entice an audience browsing the Internet (such as phrases amounting to "Click here!"). Assuming these do not contain essential information, and that they may, in fact, get in the way of training a robust model, we cleaned the dataset as much as possible.

Descriptives

As stated before, the 1,173 articles were split into 6 categories. However, this distribution is not perfectly balanced and it is mostly skewed towards real news articles. The exact distribution is displayed in Figure 1. More than a third of the articles

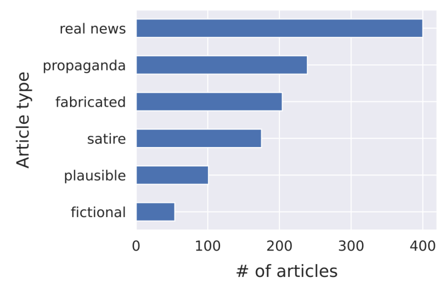

Figure 1. Article distribution by label.

belong in the real news category. Together with the plausible news, these could be considered true news informally and they would constitute over 40% of the entire dataset. The fabricated and propaganda news articles are closer to the “ideal” value, with 37%, while the fictional and satire articles constitute less than 20% of the dataset. In particular, the fictional category only contains 54 articles, being the smallest one.

Another relevant metric is the length of each article. More than half of all articles are roughly less than 500-words long. However, around 150 contain more than 1,000 words. This metric is not only relevant because it gives an idea of how many resources are required to process them, but also because it helps in choosing certain parameters of the models described later (for example, it can help us decide how long the sequences processed by BERT should be). Also, correlations may exist between the type of article and its length. In our case, it seems that fictional and satirical articles tend to be shorter than all other types, for example.

Anonymization

Because news usually talks about public figures, organizations, or institutions, such entities may become associated with certain types of articles. As such, we decided to anonymize these entities, building a different dataset to be analyzed separately. Thus, the Named Entity Recognition (NER) model from spaCy[3] was used. We anonymized all instances of people, organizations, or geopolitical entities. That being said, on manual inspection of the resulting data, it seems that common names such as “Popescu” were not always recognized. This is probably to be expected since the model has a NER accuracy of around 75-80%.

During the anonymization process, we identified approximately 17,000 references to people, 9,000 references to organizations, and 7,800 references to geopolitical entities. Among the people mentioned most often were public figures which conspiracists have connected to the COVID-19 pandemic, such as Bill and Melinda Gates, Donald Trump, or Anthony Fauci, as well as politicians from Romanian-speaking countries, such as Vlad Voiculescu and Maia Sandu. Many identified people only appear in one type of article and this could probably impact the capability of our models to generalize. Organizations that are often mentioned include the European Union and companies like Huawei or AstraZeneca, while geopolitical entities

[2] https://www.tagtog.net/fakerom/fakerom

[3] https://spacy.io/models/ro#ro_core_news_lg

referred to Romania and most of the "superpowers" of the world, such as the United States or China.

Classical Machine Learning

Before we experimented with more complex BERT-based models, we wanted to find a baseline using classical machine learning methods. Thus, we trained and evaluated models such as Support Vector Machines and Random Forests on the considered dataset. For the input features, we compared a bag-of-words representation with pre-trained word embeddings. In both cases, tokenization was done using spaCy, stopwords were removed, and the words were lemmatized.

For the bag-of-words representation, we used the TF-IDF normalization and reduced the vocabulary from approximately 26,000 tokens to just the 500 most useful ones to avoid overfitting on the small dataset. The words with the highest Chi-squared scores are presented in Table 1.

Rank	Word (Romanian and English)	Score (χ^2)
1	ziar (newspaper)	113.4
2	cip (chip)	57.1
3	roman / român (Romanian)	35.6
4	virus (virus)	19.8
5	editor (editor)	19.6
6	literă (letter)	18.4
7	doză (dose)	18.1
8	vaccinare (vaccination)	17.5

Table 1. Words with the highest predictive scores.

Pre-computed word vectors provided by the spaCy Romanian language model were considered for word embeddings. The spaCy website cites UD Romanian RRT v2.8[4] and Dumitrescu et al. [6] as sources for this model. For clarity, this framework assigns a word vector for each token in our text and then computes the mean of these embeddings to find a single 300-dimensional vector for the entire article. This vector served as the input to our models, a vector for each article. In short, the word vectors improved the performance of these models by approximately 10%, even though they were not fit specifically for our dataset (albeit the language model was trained on news articles).

BERT-based Models

We trained multiple models which used BERT as a central encoding component. We tested adding a single classification layer after BERT, as suggested by Devlin et al. [5]. We also considered an architecture based on convolutional layers, as described by Safaya et al. [14]. In all cases, the only input features were the embeddings produced by BERT. We used padding to ensure that all texts had the same number of tokens, and truncated long texts.

For most experiments, we used RoBERT [10] trained on the Romanian language and available as part of HuggingFace repository[5]. This model has multiple versions, such as RoBERT-base or RoBERT-large, depending on the size of the model. Using this Transformer model trained on Romanian text yielded better results when compared to multi-lingual models.

[4] https://github.com/UniversalDependencies/UD_Romanian-RRT
[5] https://huggingface.co/

The model adapted from the architecture proposed by Safaya et al. [14] can be seen in Figure 2. This model uses the last 4 hidden layers as features since this sometimes obtains better results than using only the last layer [7]. Apart from this, the model also uses multiple convolutional layers in parallel instead of chaining them. These layers have increasing kernel sizes, starting from 1.

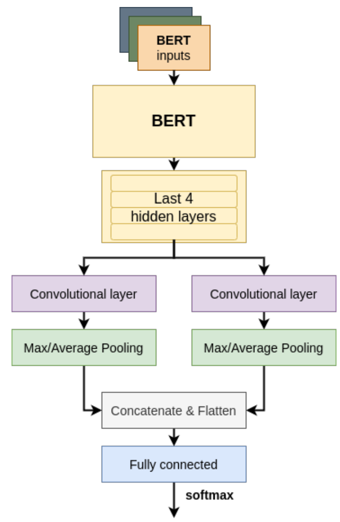

Figure 2. A BERT-based model adapted from Safaya et al. [14].

Multi-task Learning

With the goal of improving the performance of the simple BERT-based model, we experimented with multi-task learning (MTL). The multi-task learning approach refers to training a single model to solve multiple different tasks at the same time with the strive that the model will learn shared knowledge between these tasks. This can force a model to generalize better because overfitting all tasks should be more difficult. In our case, we introduced besides the original classification task two new binary classification tasks (i.e., whether the article is related to health and whether it is related to politics) and trained a single model to solve them at the same time.

As mentioned before, the articles were annotated with different subjects and topics. Many of the articles covered multiple topics, so we made a separate, binary prediction for each selected topic. Almost all articles covered the COVID subject, which we discarded since it did not provide meaningful information. The two categories to be predicted were *health* and *politics*. These were the most common ones (see Figure 3). However, even these categories were fairly rare in a sense, making the prediction task quite difficult. For example, the articles which were not tagged with the health category outnumbered those tagged by almost 3:1; so, the two classes were imbalanced.

The resulting architecture was fairly simple and can be seen in Figure 4. The BERT model is shared by all tasks with the aim to adapt well to each of them. Then, we use a different

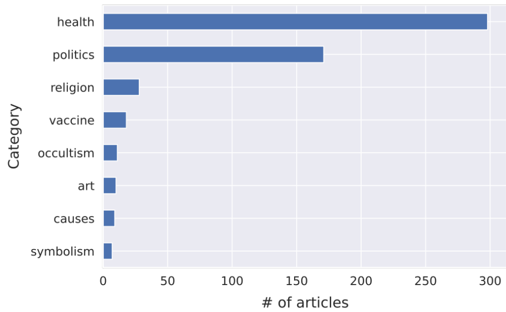

Figure 3. Distribution of the most common categories.

prediction head for each sub-task. More layers could be placed between the BERT output and the final classification layers to build a more complex representation of the problem, but given the limited number of examples, we only experimented with this simple model.

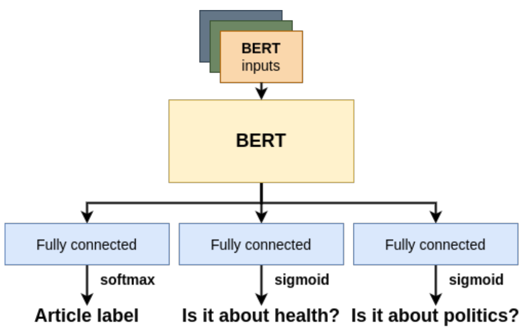

Figure 4. The architecture for Multi-task Learning.

Performance Metrics

Accuracy is one of the most common metrics used to evaluate the performance of classifiers, being equal to the fraction of correct predictions made by the model. However, in cases where there is a significant class imbalance, models can achieve high accuracy scores while performing well only on a small number of select classes. Because of this, metrics such as precision and recall provide better estimates of the true performance of a model. The F1-score is the harmonic mean of precision and recall. While we chose to present both accuracy scores and F1-scores, more attention should probably be paid to the F1-scores. The macro-F1 is computed as the arithmetic mean (i.e., unweighted mean) of all the per-class F1 scores; in contrast, the weighted F1-scores take into account the number of samples in each class. In addition, we also report confusion matrices to have a better idea of how a model performs for each class. These matrices compare quantitatively predicted labels with actual labels, enabling us to observe which classes are mislabeled most often.

RESULTS

Baseline

Out of all evaluated classical machine learning models, the SVM and XGBoost classifier based on gradient boosted trees seemed to perform the best, with accuracies and F1-scores over .60, even approaching .70. These results can be seen in Table 2 and Table 3. The confusion matrix obtained by the best SVM model is depicted in Figure 5.

Model	Accuracy	Macro F1	Weighted F1
XGBoost	**0.58**	**0.53**	**0.56**
Random Forest	**0.58**	0.44	0.51
SVM	0.53	0.39	0.48
Naive-Bayes	0.43	0.22	0.33

Table 2. Results of classical machine learning methods using bag-of-words representations.

Model	Accuracy	Macro F1	Weighted F1
XGBoost	0.65	0.60	0.63
Random Forest	0.63	0.53	0.59
SVM	**0.67**	**0.62**	**0.65**
Naive-Bayes	0.60	0.49	0.56

Table 3. Results of classical machine learning methods using word embeddings.

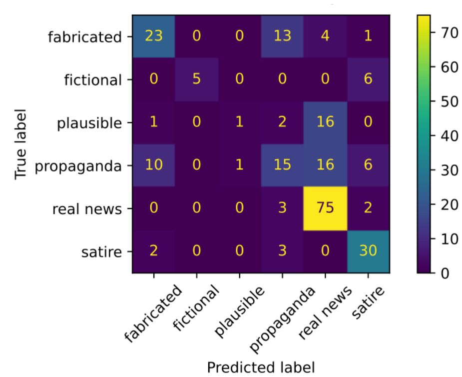

Figure 5. Confusion matrix using SVM and spaCy's word embeddings.

BERT-based Models

In general, RoBERT-base seemed to be easier to adapt to the task than RoBERT-large, obtaining better results by 1-2% (see Table 4). All the best scores were achieved using RoBERT-base.

The confusion matrix obtained by the best BERT model can be seen in Figure 6. Compared to the baseline, it seems that this model can better discern between real news and plausible news or propaganda, and it classifies satire much better, with an F1-score of .96.

The anonymized task proved to be more difficult (see Table 5). In a manner, this was expected, as the anonymized dataset hides a lot of information and, in our dataset, many of the anonymized entities were correlated with certain types of articles.

Multi-task Learning

The scores obtained on the main classification task after training the network through MTL are presented in Table 6. Once

Model	Seq. Len.	Parameters	Accuracy	Macro F1	Weighted F1
RoBERT-base + CNN	512	4 hidden-state layers	**0.75**	**0.75**	**0.75**
RoBERT-base	256	weighted classes	0.74	**0.75**	**0.75**
RoBERT-base + CNN	512	1 hidden-state layer	**0.75**	0.71	0.73
RoBERT-large	512	-	0.73	0.71	0.72

Table 4. The best results obtained by BERT-based models.

Model	Seq. Len.	Parameters	Accuracy	Macro F1	Weighted F1
RoBERT-base	256	-	**0.71**	**0.65**	**0.68**
RoBERT-base	256	weighted classes	0.70	**0.65**	**0.68**
RoBERT-large	512	-	0.68	0.62	0.64

Table 5. The results obtained by BERT+FC on the anonymized dataset.

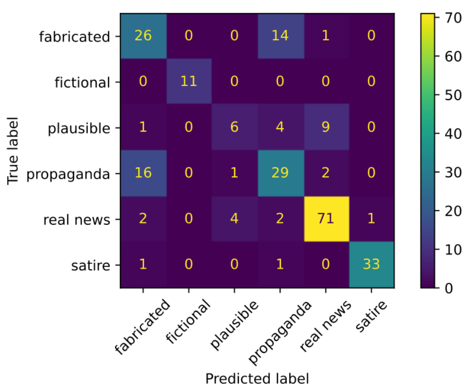

Figure 6. Confusion matrix using the best BERT model.

Model	Accuracy	Macro F1	Weighted F1
RoBERT-large	0.74	**0.73**	**0.74**
RoBERT-base	**0.75**	0.71	0.73
RoBERT-small	0.59	0.52	0.54

Table 6. The best results on the main classification task obtained by the MTL networks.

again, RoBERT-base and RoBERT-large obtained similar results, with RoBERT-large seeming a bit better in this case. At the same time, the best overall accuracy was obtained by RoBERT-base, but we need to keep in mind that the classes were imbalanced, so accuracy might be less relevant than the weighted F1-score.

DISCUSSION

It is clear from the previous results that the BERT-based models performed better than classical machine learning models. This is to be expected since we are dealing with much more complex models which were pretrained on a large amount of data. However, there were important differences even for methods such as SVM between using the bag-of-words representation and word embeddings, where we saw an improvement in F1-scores of more than 10%.

Moving to the BERT-based models, we observe that probably the main improvements consist in learning to better classify real news and discerning between satire and all other types of news. The fictional category is also well classified by most models, usually obtaining F1-scores of 1 or very close to 1.

It is probable that, since these models are more complex and benefit from BERT's understanding of language, they simply learn more characteristics which set each category apart and, thus, are capable to classify articles better. But looking at two of the categories which show the greatest improvement, we believe the length of the articles could have a small role to play in this improvement. The fictional and satire classes have the shortest articles in the dataset. While all other types of articles can have lengths from tens of words to thousands of words, these two types are usually well below the 500-word mark. As seen before, the sequence length used for BERT input was usually between 256 and 512 tokens. We considered these sizes to be appropriate since the majority of articles had fewer than 500 words and each word is only split into a few tokens at most.

Looking again at the confusion matrices from the previous section, most misclassifications happen between real and plausible news, as well as between fabricated and propaganda articles. Most models do not predict articles as being plausible too often, instead jumping straight to the conclusion that they are real, which causes a large number of false positives in the real news category. The simple model which added a single linear layer on top of BERT seemed to perform a little better in this regard, classifying a few articles as plausible, but misclassifying some of these. This model seemed to confuse plausible and real articles in both directions.

Such a situation can be seen both as a limit and as a gain. As a limit, it can be explained by the very design of the six categories of news used so far: since the difference between real news and plausible/authentic news lies simply in the fact that the former tells the truth, while the latter speaks truely [17]; however, this difference can be expressed sometimes through implicit meanings that cannot be identified by an automated analysis. At the same time, the difference between propaganda and fabricated news may more often be expressed through contextual means rather than textual ones. However, this limit can also be seen as a gain, given that the confusion between subcategories (real versus plausible and propaganda versus fabricated news) indirectly confirms the coherence and thus the legitimacy of the main overararching categories of news (i.e., true versus fake).

Regarding anonymization, the classification of the anonymized dataset could be considered almost a different task. This is because, at least in our case, anonymization removed a lot of useful information. More than ten thousand entities were replaced with generic labels, and this caused the models' performance to be constantly lower than on the regular dataset. However, it is possible that this is not characteristic of our dataset, but of news articles in general. News articles tend to condense information, in a sense, and very often they refer to important people, be they politicians, athletes, or celebrities. Most articles probably talk about a public person, an organization, or an institution, and all this data is hidden from the models in the anonymized dataset. This way, the models are forced to focus only on the plausibility of the described actions and this proved to be relatively difficult.

Usually, models which used weighted classes during training performed a bit better than those without. It is probably important to mention that there are many ways to deal with imbalanced classes, some of them being oversampling (reusing some samples from the minority classes) or assigning different numeric weights when computing certain values. In our experiments, we only used the weighted method.

As suggested by Safaya et al. [14], using more hidden states from BERT seemed to generate better results in general. However, different from the model described in the previous study, our best results were obtained when using relatively few parallel convolutional layers, 1 or 2 at most. This can probably be attributed to the small number of samples in our dataset. We also experimented with 3 or 4 layers, but Optuna (i.e., a framework used to search for hyperparameter values) [1] quickly moved away from those values as it searched for better hyperparameters.

Compared to the models with only linear layers on top of BERT, the CNN-based models seemed to become more accurate even without fine-tuning BERT. While the simpler models really became accurate only after the fine-tuning process, CNNs often reached 50-60% accuracy on the training and validation sets after 1-2 epochs of training with a frozen BERT. It is important to mention that researchers sometimes obtain good results even without fine-tuning BERT [7] and that the process of fine-tuning is not always stable, either.

As seen before, multi-task learning did not help us obtain better results in our experiments. However, the chosen tasks were probably not very well-fit for our purpose, since the number of labeled samples was small and this caused a high class imbalance.

Automated systems for fake news detection, such as those previously described, could read articles and indicate a degree of trustworthiness of their contents, thus resulting in many possible applications. For example, we developed a prototype web application (see Figure 7) that integrates the best classification model and where a user submits an article to analyze its credibility. Similarly, a browser extension could be implemented to automatically alert users to pay attention when reading news articles that have a high likelihood to propagate fake news. Another use case is to employ such systems in

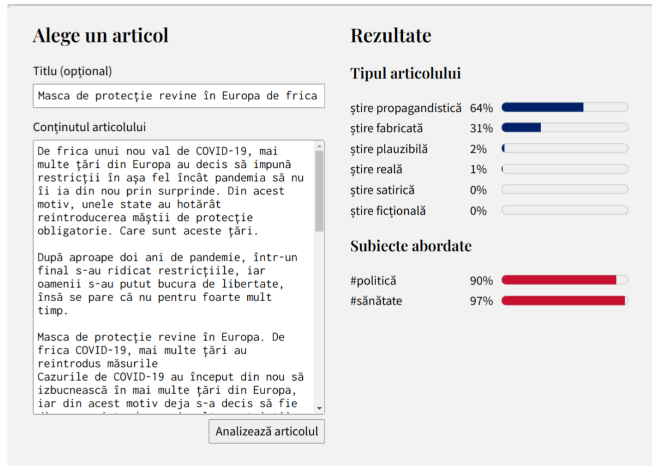

Figure 7. Our web application to detect fake news in Romanian.

social media websites to filter news articles made public or, at least, indicate if certain news pieces are suspicious.

CONCLUSIONS AND FUTURE WORK

In this project, we explored a dataset containing Romanian news articles of different types and described multiple models based on BERT or other machine learning methods for classifying these articles into fake news, true news, and more types in between. Fine-tuning BERT yielded much better results than using classical machine learning methods with pre-trained word embeddings, but other attempts at improving the performance, such as utilizing Multi-task Learning, did not succeed.

For future work, the most important step is to enrich the input data by computing certain linguistic features. This was probably the main reason why our models did not perform better. For computing these features, we could try incorporating metrics such as the difficulty of the text or the number of mentioned controversial public figures. Features derived from sentiment analysis tasks might also be useful since fake news articles tend to try to influence the reader's emotions.

Another improvement is to build a more balanced dataset or experiment with more methods of combatting this inequality. This could fall under the category of improving the dataset, which could also benefit from more samples. However, growing this dataset larger could prove to be quite difficult, since it requires the direct involvement of annotators.

In regard to the anonymization task, improvements could be made to the process of anonymization itself. In our case, spaCy was useful for finding names of people and organizations, but it was not very precise in doing so, while also removing a good amount of common, useful words in the process. Perhaps a method based on a dictionary containing lists of known organizations, institutions, or public figures, would yield better results.

If we want to continue applying multi-task and transfer learning for improving our fake news classifier, we could start by looking for better additional tasks to train together with. Once a bigger dataset is developed, this could become easier. At the same time, the original dataset contained many unlabeled samples and it is possible that they contain information that could be adapted into a good task for MTL.

Finally, a better approach to searching for hyperparameter values should be investigated. This could become easier if more computing resources become available, since more experiments could be run, or they could be run for longer. Sometimes, BERT-based models can achieve better results even when fine-tuning for a large number of epochs [7].

Acknowledgements
This work was supported by a grant from the Romanian Ministry of Education and Research, CNCS–UEFISCDI, project number PN-III-P1-1.1-TE-2019-1794, within PNCDI III.

REFERENCES

1. Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 2623–2631.
2. Jairo L. Alves, Leila Weitzel, Paulo Quaresma, Carlos E. Cardoso, and Luan Cunha. 2019. Brazilian Presidential Elections in the Era of Misinformation: A Machine Learning Approach to Analyse Fake News. In *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications*, Ingela Nyström, Yanio Hernández Heredia, and Vladimir Milián Núñez (Eds.). Springer International Publishing, Cham, 72–84.
3. João Pedro Baptista and Anabela Gradim. 2022. Online disinformation on Facebook: the spread of fake news during the Portuguese 2019 election. *Journal of Contemporary European Studies* 30, 2 (2022), 297–312. DOI:http://dx.doi.org/10.1080/14782804.2020.1843415
4. Alexandre Bovet and Hernán A Makse. 2019. Influence of fake news in Twitter during the 2016 US presidential election. *Nature communications* 10, 1 (2019), 1–14.
5. Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR* abs/1810.04805 (2018). http://arxiv.org/abs/1810.04805
6. Stefan Daniel Dumitrescu and Andrei-Marius Avram. 2019. Introducing RONEC–the Romanian Named Entity Corpus. *arXiv preprint arXiv:1909.01247* (2019).
7. Samuel Flender. 2022. What exactly happens when we fine-tune BERT? (2022). Retrieved June 26, 2022 from https://towardsdatascience.com/what-exactly-happens-when-we-fine-tune-bert-f5dc32885d76.
8. Natalie Grover. 2021. Overconfident of spotting fake news? If so, you may be more likely to fall victim. (2021). Retrieved June 26, 2022 from https://www.theguardian.com/media/2021/may/31/confident-spotting-fake-news-if-so-more-likely-fall-victim.
9. Daniel Halpern, Sebastián Valenzuela, James Katz, and Juan Pablo Miranda. 2019. From belief in conspiracy theories to trust in others: Which factors influence exposure, believing and sharing fake news. In *International conference on human-computer interaction*. Springer, 217–232.
10. Mihai Masala, Stefan Ruseti, and Mihai Dascalu. 2020. RoBERT–A Romanian BERT Model. In *Proceedings of the 28th International Conference on Computational Linguistics*. 6626–6637.
11. Daniel Patrick Moynihan. 2022. Fake News, Propaganda, and Disinformation: Learning to Critically Evaluate Media Sources: What Is Fake News? (2022). Retrieved June 26, 2022 from https://guides.library.cornell.edu/evaluate_news.
12. Umaru A Pate, Danjuma Gambo, Adamkolo Mohammed Ibrahim, and others. 2019. The impact of fake news and the emerging post-truth political era on Nigerian polity: A review of literature. *Studies in Media and Communication* 7, 1 (2019), 21–29.
13. Regina Rini. 2017. Fake news and partisan epistemology. *Kennedy Institute of Ethics Journal* 27, 2 (2017), E–43.
14. Ali Safaya, Moutasem Abdullatif, and Deniz Yuret. 2020. KUISAIL at SemEval-2020 Task 12: BERT-CNN for Offensive Speech Identification in Social Media. *CoRR* abs/2007.13184 (2020), 2054–2059. https://arxiv.org/abs/2007.13184
15. Dominik Stecuła. 2017. The real consequences of fake news. (2017). Retrieved June 26, 2022 from https://theconversation.com/the-real-consequences-of-fake-news-81179.
16. Simina-Maria Terian. 2021. What Is Fake News: A New Definition. *Revista Transilvania* 50, 11–12 (2021), 112–120.
17. Simina-Maria Terian. 2022. What Kind of News Is Fake News: A New Taxonomy. *Transylvanian Review* 31(Suppl.1) (2022), 70–84.
18. Sander Van Der Linden, Costas Panagopoulos, and Jon Roozenbeek. 2020. You are fake news: political bias in perceptions of fake news. *Media, Culture & Society* 42, 3 (2020), 460–470.
19. Tai-Li Wang. 2020. Does fake news matter to election outcomes? The case study of Taiwan's 2018 local elections. *Asian Journal for Public Opinion Research* 8, 2 (2020), 67–104.
20. Claire Wardle. 2017. Fake news. It's complicated. (2017). Retrieved June 26, 2022 from https://firstdraftnews.org/articles/fake-news-complicated.