

Detecting Toxicity in News Articles: Application to Bulgarian

Yoan Dinkov, Ivan Koychev

Faculty of Mathematics and Informatics
Sofia University
Sofia, Bulgaria

{jdinkov, koychev}@uni-sofia.bg

Preslav Nakov

Qatar Computing Research Institute
HBKU
Doha, Qatar

pnakov@hbku.edu.qa

Abstract

Online media aim for reaching ever bigger audience and for attracting ever longer attention span. This competition creates an environment that rewards sensational, fake, and toxic news. To help limit their spread and impact, we propose and develop a news toxicity detector that can recognize various types of toxic content. While previous research primarily focused on English, here we target Bulgarian. We created a new dataset by crawling a website that for five years has been collecting Bulgarian news articles that were manually categorized into eight toxicity groups. Then we trained a multi-class classifier with nine categories: eight toxic and one non-toxic. We experimented with different representations based on EIMo, BERT, and XLM, as well as with a variety of domain-specific features. Due to the small size of our dataset, we created a separate model for each feature type, and we ultimately combined these models into a meta-classifier. The evaluation results show an accuracy of 59.0% and a macro-F1 score of 39.7%, which represent sizeable improvements over the majority-class baseline (Acc=30.3%, macro-F1=5.2%).

1 Introduction

The number of online news sources has grown dramatically in recent years, and so has the amount of news that has been bombarding users on a daily basis, especially in social media. As people have limited time to spend reading news, capturing people's attention is getting ever harder. Media have to use various techniques to get their readers back such as bigger advertisement and better services.

Alternatively, it turns out that an easy way to attract people's attention is to use some toxicity in the articles, as people are intrigued by the unusual. Thus, our aim here is to try to detect articles that can harm, give false impressions, or deceive the readers. Such articles can use some of the following techniques:

- *Sensationalism*: overexposing insignificant or ordinary events by manipulating the main point of an article;
- *Fake news*: news that sound right, but totally misinterpret facts such as statistical data, locations, or dates, with the conscious aim of proving something wrong.
- *Conspiracy theories*: information that usually gives a lot of detail, but does not offer officially confirmed evidence or scientific research to back the claims that are being made. This is typically centered around political, or strange scientific phenomena.
- *Hate speech*: specifically targeting a person or a social group to brutalize them or to bully over the rate of normal conversation, to directly hurt or to manipulate them.

Given the proliferation of toxic news online, there have been many efforts to create tools and mechanisms to counteract their effect and spread. Such tools should help preserve and improve the reading integrity. Solving this problem is not a trivial task, and it requires a lot of effort by trained domain experts. Yet, there are limitations in how much it is possible to handle manually in a short period of time (and time is very critical as toxic content spreads fast). Thus, an attractive alternative is to use machine learning and natural language processing to automate the process of toxic news detection.

While most previous research has focused almost exclusively on English, here we target Bulgarian. In particular, we built a dataset for our experiments based on the knowledge base of *Media Scan*, which has catalogued and characterized many of the Bulgarian online media in the past five years. If a medium published a toxic news, this was recorded and the article, as well as the medium, got labelled accordingly. The analyzed media vary from digital newspapers, to media groups and blogs. For some articles there is detailed explanation with examples about why they were labelled like that. In some cases, the *Media Scan* website describes attempts to contact the authors of an article asking for clarification about some questionable facts that are being reported.

Here we use this information by performing multi-class classification over the *toxicity* labels: *fake news*, *sensations*, *hate speech*, *conspiracies*, *anti-democratic*, *pro-authoritarian*, *defamation*, *delusion*. Note that we allow multiple of these labels simultaneously. We further add a *non-toxic* label for articles that represent good news.

2 Related Work

The proliferation of false information has attracted a lot of research interest recently. This includes challenging the truthiness of news (Brill, 2001; Hardalov et al., 2016; Potthast et al., 2018), of news sources (Baly et al., 2018a, 2019; Dinkov et al., 2019), and of social media posts (Canini et al., 2011; Castillo et al., 2011), as well as studying credibility, influence, bias, and propaganda (Ba et al., 2016; Chen et al., 2013; Mihaylov et al., 2015a; Kulkarni et al., 2018; Baly et al., 2018a; Mihaylov et al., 2018; Barrón-Cedeño et al., 2019; Da San Martino et al., 2019; Zhang et al., 2019). Research was facilitated by shared tasks such as the SemEval 2017 and 2019 tasks on Rumor Detection (Derczynski et al., 2017; Gorrrell et al., 2019), the CLEF 2018 and 2019 Check-That! labs (Nakov et al., 2018; Elsayed et al., 2019b,a), which featured tasks on automatic identification (Atanasova et al., 2018, 2019) and verification (Barrón-Cedeño et al., 2018; Hasanain et al., 2019) of claims in political debates, the FEVER 2018 and 2019 task on Fact Extraction and VERification (Thorne et al., 2018), and the SemEval 2019 task on Fact-Checking in Community Question Answering Forums (Mihaylova et al., 2019), among others.

The interested reader can learn more about “fake news” from the overview by Shu et al. (2017), which adopted a data mining perspective and focused on social media. Another recent survey (Thorne and Vlachos, 2018) took a fact-checking perspective on “fake news” and related problems. Yet another survey was performed by Li et al. (2016), and it covered truth discovery in general. Moreover, there were two recent articles in *Science*: Lazer et al. (2018) offered a general overview and discussion on the science of “fake news”, while Vosoughi et al. (2018) focused on the proliferation of true and false news online.

The veracity of information has been studied at different levels: (i) claim (e.g., *fact-checking*), (ii) article (e.g., “*fake news*” *detection*), (iii) user (e.g., *hunting for trolls*), and (iv) medium (e.g., *source reliability estimation*). Our primary interest here is at the article-level.

2.1 Fact-Checking

At the claim-level, fact-checking and rumor detection have been primarily addressed using information extracted from social media, i.e., based on how users comment on the target claim (Canini et al., 2011; Castillo et al., 2011; Ma et al., 2016, 2017; Dungs et al., 2018; Kochkina et al., 2018). The Web has also been used as a source of information (Mukherjee and Weikum, 2015; Popat et al., 2016, 2017; Karadzhov et al., 2017b; Mihaylova et al., 2018; Baly et al., 2018b; Zlatkova et al., 2019).

In both cases, the most important information sources are *stance* (does a tweet or a news article agree or disagree with the claim?), and *source reliability* (do we trust the user who posted the tweet or the medium that published the news article?). Other important sources are linguistic expression, meta information, and temporal dynamics.

2.2 Stance Detection

Stance detection has been addressed as a task in its own right, where models have been developed based on data from the Fake News Challenge (Riedel et al., 2017; Thorne et al., 2017; Mohtarami et al., 2018; Hanselowski et al., 2018), or from SemEval-2017 Task 8 (Derczynski et al., 2017; Dungs et al., 2018). It has also been studied for other languages such as Arabic (Darwish et al., 2017b; Baly et al., 2018b; Mohtarami et al., 2019).

2.3 Source Reliability Estimation

Unlike stance detection, the problem of source reliability remains largely under-explored. In the case of social media, it concerns modeling the user¹ who posted a particular message/tweet, while in the case of the Web, it is about the trustworthiness of the source (the URL domain, the medium).

The source reliability of news media has often been estimated automatically based on the general stance of the target medium with respect to known manually fact-checked claims, without access to gold labels about the overall medium-level factuality of reporting (Mukherjee and Weikum, 2015; Popat et al., 2016, 2017, 2018). The assumption is that reliable media agree with true claims and disagree with false ones, while for unreliable media it is mostly the other way around. The trustworthiness of Web sources has also been studied from a Data Analytics perspective. For instance, Dong et al. (2015) proposed that a trustworthy source is one that contains very few false facts.

Note that estimating the reliability of a source is important not only when fact-checking a claim (Popat et al., 2017; Nguyen et al., 2018), but such reliability scores can be used as an important prior when addressing article-level factuality tasks such as “fake news” and click-bait detection (Brill, 2001; Hardalov et al., 2016; Karadzhov et al., 2017a; De Sarkar et al., 2018; Pérez-Rosas et al., 2018).

2.4 “Fake News” Detection

Most work on “fake news” detection has relied on medium-level labels, which were then assumed to hold for all articles from that source.

Horne and Adali (2017) analyzed three small datasets ranging from a couple of hundred to a few thousand articles from a couple of dozen sources, comparing (i) real news vs. (ii) “fake news” vs. (iii) satire, and found that the latter two have a lot in common across a number of dimensions. They designed a rich set of features that analyze the text of a news article, modeling its complexity, style, and psychological characteristics.

¹User modeling in social media and news community forums has focused on finding malicious users such as opinion manipulation *trolls*, paid (Mihaylov et al., 2015b) or just perceived (Mihaylov et al., 2015a; Mihaylov and Nakov, 2016; Mihaylov et al., 2018; Mihaylova et al., 2018), *sockpuppets* (Maity et al., 2017), *Internet water army* (Chen et al., 2013), and *seminar users* (Darwish et al., 2017a).

They found that “fake news” pack a lot of information in the title (as the focus is on users who do not read beyond the title), and use shorter, simpler, and repetitive content in the body (as writing fake information takes a lot of effort). Thus, they argued that the title and the body should be analyzed separately.

In follow-up work, Horne et al. (2018b) created a large-scale dataset covering 136K articles from 92 sources from `opensources.co`, which they characterize based on 130 features from seven categories: structural, sentiment, engagement, topic-dependent, complexity, bias, and morality. We use this set of features when analyzing news articles.

In yet another follow-up work, Horne et al. (2018a) trained a classifier to predict whether a given news article is coming from a reliable or from an unreliable (“fake news” or *conspiracy*)² source. Note that they assumed that all news from a given website would share the same reliability class. Such an assumption is fine for training (distant supervision), but it is problematic for testing, where manual documents-level labels are needed.

Pothast et al. (2018) used 1,627 articles from nine sources, whose factuality has been manually verified by professional journalists from BuzzFeed. They applied stylometric analysis, which was originally designed for authorship verification, to predict factuality (fake vs. real).

Rashkin et al. (2017) focused on the language used by “fake news” and compared the prevalence of several features in articles coming from trusted sources vs. hoaxes vs. satire vs. propaganda. However, their linguistic analysis and their automatic classification were at the article level and they only covered eight news media sources.

2.5 Work for Bulgarian

We are aware of only one piece of previous work for Bulgarian that targets toxicity. In particular, (Karadzhov et al., 2017a) built a fake news and click-bait detector for Bulgarian based on data from a hackaton.

While most of the above research has focused on isolated and specific task (such as trustworthiness, fake news, fact-checking), here we try to create a holistic approach by exploring several toxic and non-toxic labels simultaneously.

²We show in parentheses, the labels from `opensources.co` that are used to define a category.

Characteristic	Value
Toxic articles	221
Non-toxic articles	96
Media	164
Average title length (chars)	70.47
Average title length (words)	11.08
Average text length (chars)	3,613.70
Average text length (words)	556.83
Average text length (sentences)	31.64

Table 1: Statistics about the dataset.




Figure 1: Label distribution in the dataset.

3 Data

We used *Media Scan*³ as a source of toxicity labels for Bulgarian media Web sites. The site contains information about 700 media, and 150 of them are associated with at least one toxic article. Many of these toxic articles are removed after the respective media have been contacted and informed about the problems with these articles. Naturally, the author of *Media Scan* wanted to preserve the original form of the evaluated article, and thus had a link to a PDF copy in case the original HTML page was not accessible. We only crawled the HTML for articles that have not been changed or removed at the time we created the dataset.

For each Web page with a toxic label, we ran a mechanical crawler to obtain its contents. This was not very reliable as each individual medium site has its own structure, while the crawler expected more or less semantic and valid HTML to be able to process it. Thus, we manually verified the data, fixed any issues we could find and added any missing information. We ended up with a little over 200 articles with some kind of toxicity.

³<http://mediascan.gadjokov.com>

In addition to this dataset of only toxic articles, we added some “non-toxic” articles, fetched from media without toxicity examples in Media Scan: we added a total of 96 articles from 25 media.

Table 1 shows statistics about the dataset, and Figure 1 shows the distribution of the labels.

4 Method

We used a feature-rich classifier based on logistic regression and a neural network.

For each article, we extracted its title and its body. We further extracted some meta information about the corresponding news medium. As some NLP resources are only available or are better for English, we translated the articles to English, by using Google Translate API, so that we can extract features from them as explained in subsections 4.2, 4.6, 4.7, and 4.8.

4.1 LSA

We trained a Latent Semantic Analysis (LSA) model on our data. We first built TF.IDF vectors for the title and the body. Then, we applied singular value decomposition (SVD) to generate vectors of 15 dimensions for the titles and of 200 dimensions for the article bodies.

4.2 BERT

We used BERT (Devlin et al., 2019) for sentence representation, which has achieved very strong results on eleven natural language processing tasks including GLUE, MultiNLI, and SQuAD. Since then, it was used to improve over the state of the art for a number of NLP tasks. The original model was trained on English Wikipedia articles (2500M words). Due to the model complexity and to its size, it is hard to find enough data that represents a specific domain for a specific language.

We used BERT-as-a-service, which generates a vector of 768 numerical values for a given text. In its original form, this is a sentence representation tool, but we used it to generate text over the first 512 tokens of our article’s title or text. We used the multilingual cased pretrained model⁴. We experimented with all possible pooling strategies for representing the article title and its body, and we eventually chose the following pooling strategies: *REDUCE_MAX* for the title and *CLS_TOKEN* for the text of the article.

⁴<http://github.com/google-research/bert#pre-trained-models>

4.3 Stylometric Features

For the title and the body of each article, we calculate the following features:

- *avg_word_length_title*: average length of the word in the article title;
- *avg_word_length_text*: average length of the words in the article body
- *word_count_title*: number of words in the article title;
- *word_count_text*: number of words in the article body;
- *char_count_title*: number of characters in the article title;
- *char_count_text*: number of characters in the article body;
- *spec_char_count_title*: number of specific (non-alpha-numeric) characters in the article title;
- *spec_char_count_text*: number of specific (non-alpha-numeric) characters in the article body;
- *upper_char_count_title*: number of uppercase characters in the article title;
- *upper_char_count_text*: number of uppercase characters in the article body;
- *upper_word_count_title*: number of words starting with an uppercase character in the article title;
- *upper_word_count_text*: number of words starting with an uppercase character in the article body;
- *sentence_count_text*: number of sentences in the article;
- *avg_sentence_length_char_text*: average length of the sentences in the article body, in terms of characters;
- *avg_sentence_length_word_text*: average length of the sentences in the article body, in terms of words;

4.4 Media Features

We further extracted binary and numerical features characterizing the medium the article came from:

- *editor*: its value is 1 if the target medium has a designated chief editor, and it is 0 otherwise;
- *responsible_person*: its value is 1 if the target medium has a responsible person, and it is 0 otherwise;
- *bg_server*: its value is 1 if the target medium's location is in Bulgaria, and it is 0 otherwise;
- *popularity*: reciprocal value of the target medium's rank the Web traffic analysis platform Alexa⁵;
- *domain_person*: its value is 1 if the target medium has a designated owner, and it is 0 otherwise;
- *days_existing*: number of days between when the medium was created and 01.01.2019. As this value is quite large, we take the logarithm thereof. For example, a medium created on January 1, 2005 would have 5,113 days of existence, which would correspond to 3.70 as this feature's value.

4.5 XLM

We further used cross-lingual representations from the Facebook's XLM model (Lample and Conneau, 2019), which creates cross-lingual representations based on the Transformer, similarly to BERT. This model is pretrained for 15 languages including Bulgarian and English. We use their pretrained models, which were fine-tuned for Cross-lingual Natural Language Interence (XNLI) tasks. This yielded a 1024-demnsional representation for the title, and another one for the article body.

4.6 Universal Sentence Encoder

We also extracted representation using Google's Universal Sentence Encoder, or USE, (Cer et al., 2018). We used the pretrained model from TF Hub⁶. As the model is only available for English, we used the translations of the news articles. We passed the model the first 300 tokens for each title or body to generate 512-dimensional vectors.

⁵<http://www.alexa.com/>

⁶<http://tfhub.dev/google/universal-sentence-encoder/2>

4.7 EIMo

Next, we use deep contextualized word representations from EIMo, which uses generative bidirectional language model pre-training (Peters et al., 2018). The model yields 1024-dimensional representation, which we generate separately for the article title and for its body.

4.8 NELA Features

Finally, we use features from the NELA toolkit (Horne et al., 2018a), which were previously shown useful for detecting fake news, political bias, etc. The toolkit implements 129 features, which we extract separately for the article title and for its body:

- **Structure:** POS tags, linguistic features (function words, pronouns, etc.), and features for clickbait title classification from (Chakraborty et al., 2016);
- **Sentiment:** sentiment scores using lexicons (Recasens et al., 2013; Mitchell et al., 2013) and full systems (Hutto and Gilbert, 2014);
- **Topic:** lexicon features to differentiate between science topics and personal concerns;
- **Complexity:** type-token ratio, readability, number of cognitive process words (identifying discrepancy, insight, certainty, etc.);
- **Bias:** features modeling bias (Recasens et al., 2013; Mukherjee and Weikum, 2015) and subjectivity (Horne et al., 2017);
- **Morality:** features based on the Moral Foundation Theory (Graham et al., 2009) and lexicons (Lin et al., 2018)

A summary of all features is shown in Table 2.

Feature Group	Title	Body
BERT	768	768
EIMo	1,024	1,024
LSA	15	200
NELA	129	129
Stylometry	19	6
USE	512	512
XLM	1,024	1,024
Media	6	

Table 2: Summary of our features.

5 Experiments and Evaluation

5.1 Experimental Setup

We used logistic regression as our main classification method. As we have a small dataset, we performed 5-fold cross-validation. For evaluation, we used accuracy and macro-average F_1 score. The results are presented in Table 3.

Our baseline approach is based on selecting the most frequent class, i.e., *non-toxic*, which covers 30.30% of the dataset (see Table 1).

5.2 Individual Models

We evaluated a total of 14 setups for feature combination. Four of them represent features generated from the original article’s title and body as well as a combination thereof. The next four setups present feature sets generated from the English translation as well as a combination thereof. The final section of Table 3 shows three setups that are somewhat language-independent: meta media, all features combined together as well as a meta classifier. We tuned the logistic regression for each individual experimental setup, using an additional internal cross-validation for the training part of each experiment in the 5-fold cross-validation. In total, 15,000 additional experiments have been conducted to complete the fine-tuning.

We can see in Table 3 that the BERT features (setups 2, 9) perform well both for English and for Bulgarian. The feature combinations (setups 6, 11, 13) do not yield good results as this increases the number of features, while the number of training examples remains limited. Using only the right 6 meta features about the target medium yields 12% improvement over the baseline. Interestingly, LSA turns out to be the best text representation model.

5.3 Meta Classifier

Next we tried a meta classifier. For this purpose, we extracted the posterior probabilities of the individual classifiers in Table 3 (2-5, 7-10, 12). Then, we trained a logistic classifier on these posteriors (we made sure that we do not leak information about the labels when training the meta classifier).

We can see in Table 3 that the meta classifier yielded the best results, outperforming all of the individual models, and achieving 3.5% absolute gain in terms of accuracy.

Language	N	Feature Set	Dimension	Accuracy	F1-macro
-	1	Baseline	-	30.30	05.17
BG	2	BERT(title), BERT(text)	1,536	47.69	32.58
	3	XLM(title), XLM(text)	2,048	38.50	24.58
	4	Styl(title), Styl(text)	15	31.89	08.51
	5	LSA(title), LSA(text)	215	55.59	42.11
	6	Bulgarian combined	3,824	39.43	24.38
EN	7	USE(title), USE(text)	1,024	53.70	40.68
	8	NELA(title), NELA(text)	258	36.36	23.04
	9	BERT(title), BERT(text)	1,536	52.05	39.78
	10	EIMO(title), EIMO(text)	2,048	54.60	40.95
	11	English combined	4,878	45.45	31.42
-	12	Media meta	6	42.04	15.64
-	13	All combined	8,694	38.16	26.04
-	14	Meta classifier	153	59.06	39.70

Table 3: Evaluation results.

True labels \ Predicted labels	anti-democratic	delusion	defamation	conspiracy	non-toxic	pro-authoritarian	hate speech	sensation	fake news
anti-democratic	0	0	3	0	2	1	2	0	7
delusion	0	6	1	3	4	0	0	0	6
defamation	0	0	20	0	6	0	2	0	3
conspiracy	0	1	2	34	3	0	0	0	9
non-toxic	0	0	1	2	87	0	1	0	5
pro-authoritarian	0	0	3	4	2	0	0	0	4
hate speech	0	0	1	2	4	0	16	0	3
sensation	0	0	1	0	9	0	2	1	4
fake news	0	2	1	14	7	0	3	0	23

Figure 2: Confusion matrix.

5.4 Analysis

Figure 2 shows a confusion matrix over the entire dataset for the best model in Table 3, namely experiment 14. We can see that the model works best for the biggest *non-toxic* class. A decent chunk of *fake news* samples are misclassified as *conspiracy* as those two classes are the second and the third largest ones. For three of the labels, there are hardly any predictions; these are the smallest classes, and three of them combined cover less than 18% of the dataset.




Figure 3: Our neural network.

5.5 What Did Not Work

Oversampling We evaluated all models from Table 3 but with oversampling the small classes. We tried simple random oversampling as well as the more complex SMOTE (Chawla et al., 2002) model. None of these techniques yielded improvements.

Neural Network We also tried a feed forward neural network with two hidden densely connected layers, with 64 nodes (ReLU activation) and 32 nodes (Tanh activation), and a dropout rate of 0.35 for each of them; see Figure 3. We used Adam for optimization, and we tried various parameter values, but we could not get improvements, possibly due to the small size of our dataset.

6 Discussion

Below we compare the performance of the models when working with Bulgarian vs. English resources, and we further discuss issues related to the small size of our dataset.

6.1 Language-Related Issues

The first part of the feature comparison in Table 3 is between English and Bulgarian, and it is interesting to compare them. The first comparison is between the BERT features. Even though we used a pre-trained BERT model, with same pooling techniques there is close to 4.5% absolute improvement when using the English translation. This is probably due to the English BERT being trained on more data as the English Wikipedia is much bigger for English: 5.7M English articles vs. just 250K Bulgarian articles.

Another notable comparison is between the types of models. Two of the Bulgarian feature sets are created via local models (experiments 4 and 5), while all of the English experiments are from transfer-learning. We can see that LSA (experiment 5) is the best feature set, and one can argue that this is to be expected. On such a small dataset in a non-English language, it is hard to represent the text with pre-trained models. Nevertheless, we can see that a combination between only pre-trained models (experiment 11) performs better compared to fusion between local and transfer-learning models (experiment 6).

6.2 Data Size Issues

There are several aspects of the above experiments where we can observe the negative effect of having insufficient data. First, in experiments 6, 11, 12 in Table 3, we can see that the combination of features performs worse in each language group compared to single feature types.

Another place where we felt we had insufficient data was in the neural network experiments, where we had many more parameters to train than in the simple logistic regression.

A related problem is that of class imbalance: we have seen in Section 5.4 above that the smallest classes were effectively ignored even by our best classifier. We can see in Figure 1 that those three classes have less than 60 articles combined, while the “non-toxic” only had 96 articles.

7 Conclusion and Future Work

We have presented experiments in detecting the toxicity of news articles. While previous research was mostly limited to English, here we focused on Bulgarian. We created a new dataset by crawling a website that has been collecting Bulgarian news articles and manually categorized them in eight toxicity groups. We then trained a multi-class classifier with nine categories: eight toxic and one non-toxic. We experimented with a variety of representations based on EIMo, BERT, and XLM. We further used a variety of domain-specific features, which we eventually combined in a meta classifier. The evaluation results show an accuracy of 59.0% and a macro-F1 score of 39.7%, which represent sizable improvements over the majority-class baseline (Acc=30.3%, macro-F1=5.2%).

In future work, we plan to extend and also to balance the dataset. This can be achieved by either exploring another source for articles using the methodology of *Media Scan*, or by processing the unstructured PDF article, which we ignored in the present study. We also plan to explore new information sources. From a technical point of view, we would like to improve the neural network architecture as well as the oversampling techniques (with possible combination with undersampling).

Data and Code We are releasing all of the code for our experiments in a public repository that can be found in GitHub⁷ with explanations about how to reproduce our environment. In that repository, we further release the full dataset together with the generated features, all the textual data, all the translations and all the meta data about the articles.

Acknowledgements

This research is part of the Tanbih project,⁸ which aims to limit the effect of “fake news”, propaganda and media bias by making users aware of what they are reading. The project is developed in collaboration between the Qatar Computing Research Institute (QCRI), HBKU and the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL).

This research is also partially supported by Project UNITE BG05M2OP001-1.001-0004 funded by the OP “Science and Education for Smart Growth” and the EU via the ESI Funds.

⁷github.com/yoandinkov/ranlp-2019

⁸<http://tanbih.qcri.org/>

References

- Pepa Atanasova, Lluís Màrquez, Alberto Barrón-Cedeño, Tamer Elsayed, Reem Suwaileh, Wajdi Zaghoulani, Spas Kyuchukov, Giovanni Da San Martino, and Preslav Nakov. 2018. Overview of the CLEF-2018 CheckThat! lab on automatic identification and verification of political claims, Task 1: Check-worthiness. In *CLEF 2018 Working Notes*. CEUR-WS.org, Avignon, France.
- Pepa Atanasova, Preslav Nakov, Georgi Karadzhov, Mitra Mohtarami, and Giovanni Da San Martino. 2019. Overview of the CLEF-2019 CheckThat! Lab on Automatic Identification and Verification of Claims. Task 1: Check-Worthiness. In *CLEF 2019 Working Notes*. CEUR-WS.org, Lugano, Switzerland.
- Mouhamadou Lamine Ba, Laure Berti-Equille, Kushal Shah, and Hossam M. Hammady. 2016. VERA: A platform for veracity estimation over web data. In *Proceedings of the 25th International Conference Companion on World Wide Web*. Montréal, Québec, Canada, WWW '16, pages 159–162.
- Ramy Baly, Georgi Karadzhov, Dimitar Alexandrov, James Glass, and Preslav Nakov. 2018a. Predicting factuality of reporting and bias of news media sources. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium, EMNLP '18, pages 3528–3539.
- Ramy Baly, Georgi Karadzhov, Abdelrhman Saleh, James Glass, and Preslav Nakov. 2019. Multi-task ordinal regression for jointly predicting the trustworthiness and the leading political ideology of news media. In *Proceedings of the 17th Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Minneapolis, MN, USA, NAACL-HLT '19, pages 2109–2116.
- Ramy Baly, Mitra Mohtarami, James Glass, Lluís Màrquez, Alessandro Moschitti, and Preslav Nakov. 2018b. Integrating stance detection and fact checking in a unified corpus. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*. New Orleans, LA, USA, NAACL-HLT '18, pages 21–27.
- Alberto Barrón-Cedeño, Giovanni Da San Martino, Israa Jaradat, and Preslav Nakov. 2019. Propopy: Organizing the news based on their propagandistic content. *Information Processing & Management* 56(5):1849 – 1864.
- Alberto Barrón-Cedeño, Tamer Elsayed, Reem Suwaileh, Lluís Màrquez, Pepa Atanasova, Wajdi Zaghoulani, Spas Kyuchukov, Giovanni Da San Martino, and Preslav Nakov. 2018. Overview of the CLEF-2018 CheckThat! lab on automatic identification and verification of political claims, Task 2: Factuality. In *CLEF 2018 Working Notes*. CEUR-WS.org, Avignon, France.
- Ann M Brill. 2001. Online journalists embrace new marketing function. *Newspaper Research Journal* 22(2):28–40.
- Kevin R. Canini, Bongwon Suh, and Peter L. Pirolli. 2011. Finding credible information sources in social networks based on content and social structure. In *Proceedings of the IEEE International Conference on Privacy, Security, Risk, and Trust, and the IEEE International Conference on Social Computing*. Boston, MA, USA, SocialCom/PASSAT '11, pages 1–8.
- Carlos Castillo, Marcelo Mendoza, and Barbara Poblete. 2011. Information credibility on Twitter. In *Proceedings of the 20th International Conference on World Wide Web*. Hyderabad, India, WWW '11, pages 675–684.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Abhijnan Chakraborty, Bhargavi Paranjape, Kakarla Kakarla, and Niloy Ganguly. 2016. Stop clickbait: Detecting and preventing clickbaits in online news media. In *Proceedings of the 2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. San Francisco, CA, USA, ASONAM '16, pages 9–16.
- Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. 2002. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research* 16:321–357.
- Cheng Chen, Kui Wu, Venkatesh Srinivasan, and Xudong Zhang. 2013. Battling the Internet Water Army: detection of hidden paid posters. In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. Niagara, Canada, ASONAM '13, pages 116–120.
- Giovanni Da San Martino, Seunghak Yu, Alberto Barron-Cedeno, Rostislav Petrov, and Preslav Nakov. 2019. Fine-grained analysis of propaganda in news articles. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Hong Kong, China, EMNLP '19.
- Kareem Darwish, Dimitar Alexandrov, Preslav Nakov, and Yelena Mejova. 2017a. Seminar users in the Arabic Twitter sphere. In *Proceedings of the 9th International Conference on Social Informatics*. Oxford, UK, SocInfo '17, pages 91–108.
- Kareem Darwish, Walid Magdy, and Tahar Zanouda. 2017b. Improved stance prediction in a user similarity feature space. In *Proceedings of the Conference on Advances in Social Networks Analysis and Mining*. Sydney, Australia, ASONAM '17, pages 145–148.

- Sohan De Sarkar, Fan Yang, and Arjun Mukherjee. 2018. Attending sentences to detect satirical fake news. In *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, NM, USA, COLING '18, pages 3371–3380.
- Leon Derczynski, Kalina Bontcheva, Maria Liakata, Rob Procter, Geraldine Wong Sak Hoi, and Arkaitz Zubiaga. 2017. SemEval-2017 Task 8: RumourEval: Determining rumour veracity and support for rumours. In *Proceedings of the 11th International Workshop on Semantic Evaluation*. Vancouver, Canada, SemEval '17, pages 60–67.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. 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*. Minneapolis, MN, USA, NAACL-HLT '2019, pages 4171–4186.
- Yoan Dinkov, Ahmed Ali, Ivan Koychev, and Preslav Nakov. 2019. Predicting the leading political ideology of Youtube channels using acoustic, textual and metadata information. In *Proceedings of the 20th Annual Conference of the International Speech Communication Association*. Graz, Austria, INTERSPEECH '19.
- Xin Luna Dong, Evgeniy Gabilovich, Kevin Murphy, Van Dang, Wilko Horn, Camillo Lugaresi, Shao-hua Sun, and Wei Zhang. 2015. Knowledge-based trust: Estimating the trustworthiness of web sources. *Proc. VLDB Endow.* 8(9):938–949.
- Sebastian Dungs, Ahmet Aker, Norbert Fuhr, and Kalina Bontcheva. 2018. Can rumour stance alone predict veracity? In *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, NM, USA, COLING '18, pages 3360–3370.
- Tamer Elsayed, Preslav Nakov, Alberto Barrón-Cedeño, Maram Hasanain, Reem Suwaileh, Pepa Atanasova, and Giovanni Da San Martino. 2019a. CheckThat! at CLEF 2019: Automatic identification and verification of claims. In *Proceedings of the 41st European Conference on Information Retrieval*. Cologne, Germany, ECIR '19, pages 309–315.
- Tamer Elsayed, Preslav Nakov, Alberto Barrón-Cedeño, Maram Hasanain, Reem Suwaileh, Giovanni Da San Martino, and Pepa Atanasova. 2019b. Overview of the CLEF-2019 CheckThat!: Automatic identification and verification of claims. In *Experimental IR Meets Multilinguality, Multimodality, and Interaction*. Springer, Lugano, Switzerland.
- Genevieve Gorrell, Elena Kochkina, Maria Liakata, Ahmet Aker, Arkaitz Zubiaga, Kalina Bontcheva, and Leon Derczynski. 2019. SemEval-2019 task 7: RumourEval, determining rumour veracity and support for rumours. In *Proceedings of the 13th International Workshop on Semantic Evaluation*. Minneapolis, MN, USA, SemEval '19, pages 845–854.
- Jesse Graham, Jonathan Haidt, and Brian A Nosek. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology* 96(5):1029.
- Andreas Hanselowski, Avinesh PVS, Benjamin Schiller, Felix Caspelherr, Debanjan Chaudhuri, Christian M. Meyer, and Iryna Gurevych. 2018. A retrospective analysis of the fake news challenge stance-detection task. In *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, NM, USA, COLING '18, pages 1859–1874.
- Momchil Hardalov, Ivan Koychev, and Preslav Nakov. 2016. In search of credible news. In *Proceedings of the 17th International Conference on Artificial Intelligence: Methodology, Systems, and Applications*. Varna, Bulgaria, AIMS '16, pages 172–180.
- Maram Hasanain, Reem Suwaileh, Tamer Elsayed, Alberto Barrón-Cedeño, and Preslav Nakov. 2019. Overview of the CLEF-2019 CheckThat! Lab on Automatic Identification and Verification of Claims. Task 2: Evidence and Factuality. In *CLEF 2019 Working Notes*. CEUR-WS.org, Lugano, Switzerland.
- Benjamin Horne and Sibel Adali. 2017. This just in: Fake news packs a lot in title, uses simpler, repetitive content in text body, more similar to satire than real news. *CoRR* abs/1703.09398.
- Benjamin Horne, Sibel Adali, and Sujoy Sikdar. 2017. Identifying the social signals that drive online discussions: A case study of Reddit communities. In *Proceedings of the 26th IEEE International Conference on Computer Communication and Networks*. Vancouver, Canada, ICCCN '17, pages 1–9.
- Benjamin D. Horne, William Dron, Sara Khedr, and Sibel Adali. 2018a. Assessing the news landscape: A multi-module toolkit for evaluating the credibility of news. In *Proceedings of the The Web Conference*. Lyon, France, WWW '18, pages 235–238.
- Benjamin D. Horne, Sara Khedr, and Sibel Adali. 2018b. Sampling the news producers: A large news and feature data set for the study of the complex media landscape. In *Proceedings of the Twelfth International Conference on Web and Social Media*. Stanford, CA, USA, ICWSM '18, pages 518–527.
- Clayton Hutto and Eric Gilbert. 2014. VADER: A parsimonious rule-based model for sentiment analysis of social media text. In *Proceedings of the 8th International Conference on Weblogs and Social Media*. Ann Arbor, MI, USA, ICWSM '14.
- Georgi Karadzhov, Pepa Gencheva, Preslav Nakov, and Ivan Koychev. 2017a. We built a fake news & click-bait filter: What happened next will blow your mind! In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*. Varna, Bulgaria, RANLP '17, pages 334–343.

- Georgi Karadzhov, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, and Ivan Koychev. 2017b. Fully automated fact checking using external sources. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing*. Varna, Bulgaria, RANLP '17, pages 344–353.
- Elena Kochkina, Maria Liakata, and Arkaitz Zubiaga. 2018. All-in-one: Multi-task learning for rumour verification. In *Proceedings of the International Conference on Computational Linguistics*. Santa Fe, NM, USA, COLING '18, pages 3402–3413.
- Vivek Kulkarni, Junting Ye, Steve Skiena, and William Yang Wang. 2018. Multi-view models for political ideology detection of news articles. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium, EMNLP '18, pages 3518–3527.
- Guillaume Lample and Alexis Conneau. 2019. Cross-lingual language model pretraining. *CoRR* abs/1901.07291.
- David M.J. Lazer, Matthew A. Baum, Yochai Benkler, Adam J. Berinsky, Kelly M. Greenhill, Filippo Menczer, Miriam J. Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild, Michael Schudson, Steven A. Sloman, Cass R. Sunstein, Emily A. Thorson, Duncan J. Watts, and Jonathan L. Zittrain. 2018. The science of fake news. *Science* 359(6380):1094–1096.
- Yaliang Li, Jing Gao, Chuishi Meng, Qi Li, Lu Su, Bo Zhao, Wei Fan, and Jiawei Han. 2016. A survey on truth discovery. *SIGKDD Explor. Newsl.* 17(2):1–16.
- Y. Lin, J. Hoover, G. Portillo-Wightman, C. Park, M. Deghani, and H. Ji. 2018. Acquiring background knowledge to improve moral value prediction. In *Proceedings of the 2018 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*. Los Alamitos, CA, USA, ASONAM '18, pages 552–559.
- Jing Ma, Wei Gao, Prasenjit Mitra, Sejeong Kwon, Bernard J. Jansen, Kam-Fai Wong, and Meeyoung Cha. 2016. Detecting rumors from microblogs with recurrent neural networks. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*. New York, NY, USA, IJCAI '16, pages 3818–3824.
- Jing Ma, Wei Gao, and Kam-Fai Wong. 2017. Detect rumors in microblog posts using propagation structure via kernel learning. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*. Vancouver, Canada, ACL '17, pages 708–717.
- Suman Kalyan Maity, Aishik Chakraborty, Pawan Goyal, and Animesh Mukherjee. 2017. Detection of sockpuppets in social media. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work and Social Computing*. Portland, OR, USA, CSCW '17, pages 243–246.
- Todor Mihaylov, Georgi Georgiev, and Preslav Nakov. 2015a. Finding opinion manipulation trolls in news community forums. In *Proceedings of the Nineteenth Conference on Computational Natural Language Learning*. Beijing, China, CoNLL '15, pages 310–314.
- Todor Mihaylov, Ivan Koychev, Georgi Georgiev, and Preslav Nakov. 2015b. Exposing paid opinion manipulation trolls. In *Proceedings of the International Conference Recent Advances in Natural Language Processing*. Hissar, Bulgaria, RANLP '15, pages 443–450.
- Todor Mihaylov, Tsvetomila Mihaylova, Preslav Nakov, Lluís Màrquez, Georgi Georgiev, and Ivan Koychev. 2018. The dark side of news community forums: Opinion manipulation trolls. *Internet Research* 28(5):1292–1312.
- Todor Mihaylov and Preslav Nakov. 2016. Hunting for troll comments in news community forums. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*. Berlin, Germany, ACL '16, pages 399–405.
- Tsvetomila Mihaylova, Georgi Karadzhov, Pepa Atanasova, Ramy Baly, Mitra Mohtarami, and Preslav Nakov. 2019. SemEval-2019 task 8: Fact checking in community question answering forums. In *Proceedings of the 13th International Workshop on Semantic Evaluation*. Minneapolis, MN, USA, SemEval '19, pages 860–869.
- Tsvetomila Mihaylova, Preslav Nakov, Lluís Màrquez, Alberto Barrón-Cedeño, Mitra Mohtarami, Georgi Karadzhov, and James Glass. 2018. Fact checking in community forums. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*. New Orleans, LA, USA, AAAI '18, pages 879–886.
- Lewis Mitchell, Morgan R Frank, Kameron Decker Harris, Peter Sheridan Dodds, and Christopher M Danforth. 2013. The geography of happiness: Connecting Twitter sentiment and expression, demographics, and objective characteristics of place. *PloS one* 8(5):e64417.
- Mitra Mohtarami, Ramy Baly, James Glass, Preslav Nakov, Lluís Màrquez, and Alessandro Moschitti. 2018. Automatic stance detection using end-to-end memory networks. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics*. New Orleans, LA, USA, NAACL-HLT '18, pages 767–776.
- Mitra Mohtarami, James Glass, and Preslav Nakov. 2019. Contrastive language adaptation for cross-lingual stance detection. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Hong Kong, China, EMNLP '19.

- Subhabrata Mukherjee and Gerhard Weikum. 2015. Leveraging joint interactions for credibility analysis in news communities. In *Proceedings of the 24th ACM International on Conference on Information and Knowledge Management*. Melbourne, Australia, CIKM '15, pages 353–362.
- Preslav Nakov, Alberto Barrón-Cedeño, Tamer Elsayed, Reem Suwaileh, Lluís Màrquez, Wajdi Zaghouani, Pepa Atanasova, Spas Kyuchukov, and Giovanni Da San Martino. 2018. Overview of the CLEF-2018 CheckThat! lab on automatic identification and verification of political claims. In *Proceedings of CLEF*. Avignon, France, pages 372–387.
- An T. Nguyen, Aditya Kharosekar, Matthew Lease, and Byron C. Wallace. 2018. An interpretable joint graphical model for fact-checking from crowds. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*. New Orleans, LA, USA, AAAI '18, pages 1511–1518.
- Verónica Pérez-Rosas, Bennett Kleinberg, Alexandra Lefevre, and Rada Mihalcea. 2018. Automatic detection of fake news. In *Proceedings of the 27th International Conference on Computational Linguistics*. Santa Fe, NM, USA, COLING '18, pages 3391–3401.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*. New Orleans, LA, USA, NAACL-HLT '18, pages 2227–2237.
- Kashyap Popat, Subhabrata Mukherjee, Jannik Strötgen, and Gerhard Weikum. 2016. Credibility assessment of textual claims on the web. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. Indianapolis, IN, USA, CIKM '16, pages 2173–2178.
- Kashyap Popat, Subhabrata Mukherjee, Jannik Strötgen, and Gerhard Weikum. 2017. Where the truth lies: Explaining the credibility of emerging claims on the Web and social media. In *Proceedings of the 26th International Conference on World Wide Web Companion*. Perth, Australia, WWW '17, pages 1003–1012.
- Kashyap Popat, Subhabrata Mukherjee, Jannik Strötgen, and Gerhard Weikum. 2018. CredEye: A credibility lens for analyzing and explaining misinformation. In *Proceedings of The Web Conference 2018*. Lyon, France, WWW '18, pages 155–158.
- Martin Potthast, Johannes Kiesel, Kevin Reinartz, Janek Bevendorff, and Benno Stein. 2018. A stylistic inquiry into hyperpartisan and fake news. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*. Melbourne, Australia, ACL '18, pages 231–240.
- Hannah Rashkin, Eunsol Choi, Jin Yea Jang, Svitlana Volkova, and Yejin Choi. 2017. Truth of varying shades: Analyzing language in fake news and political fact-checking. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Copenhagen, Denmark, EMNLP '17, pages 2931–2937.
- Marta Recasens, Cristian Danescu-Niculescu-Mizil, and Dan Jurafsky. 2013. Linguistic models for analyzing and detecting biased language. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*. Sofia, Bulgaria, ACL '13, pages 1650–1659.
- Benjamin Riedel, Isabelle Augenstein, Georgios P Spithourakis, and Sebastian Riedel. 2017. A simple but tough-to-beat baseline for the Fake News Challenge stance detection task. *ArXiv:1707.03264*.
- Kai Shu, Amy Sliva, Suhang Wang, Jiliang Tang, and Huan Liu. 2017. Fake news detection on social media: A data mining perspective. *SIGKDD Explor. Newsl.* 19(1):22–36.
- James Thorne, Mingjie Chen, Giorgos Myrianthous, Jiashu Pu, Xiaoxuan Wang, and Andreas Vlachos. 2017. Fake news stance detection using stacked ensemble of classifiers. In *Proceedings of the EMNLP Workshop on Natural Language Processing meets Journalism*. Copenhagen, Denmark, pages 80–83.
- James Thorne and Andreas Vlachos. 2018. Automated fact checking: Task formulations, methods and future directions. In *Proceedings of the International Conference on Computational Linguistics*. Santa Fe, NM, USA, COLING '18, pages 3346–3359.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: a large-scale dataset for fact extraction and VERification. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*. New Orleans, LA, USA, NAACL-HLT '18, pages 809–819.
- Soroush Vosoughi, Deb Roy, and Sinan Aral. 2018. The spread of true and false news online. *Science* 359(6380):1146–1151.
- Yifan Zhang, Giovanni Da San Martino, Alberto Barrón-Cedeño, Salvatore Romeo, Jisun An, Hae-woon Kwak, Todor Staykovski, Israa Jaradat, Georgi Karadzhov, Ramy Baly, Kareem Darwish, and Preslav Nakov James Glass. 2019. Tanbih: Get to know what you are reading. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Hong Kong, China, EMNLP '19.
- Dimitrina Zlatkova, Preslav Nakov, and Ivan Koychev. 2019. Fact-checking meets fauxtography: Verifying claims about images. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Hong Kong, China, EMNLP '19.