Discussion Manipulation, Language and Domain Dependent Models: An Overview

Filip Chalás*, Igor Stupavský†, and Valentino Vranić‡
Institute of Informatics, Information Systems, and Software Engineering
Faculty of Informatics and Information Technologies
Slovak University of Technology in Bratislava
*xchalas@stuba.sk, ORCID: 0000-0002-8652-8417
†igor.stupavsky@stuba.sk, ORCID: 0000-0002-0877-433X
‡vranic@stuba.sk, ORCID: 0000-0001-9044-4593

***Abstract*—The amount of antisocial online behavior (AOB) in the political field has been on the rise in recent years. In the research, we are dealing with the reason for confusing AOB with other forms of antisocial behavior. We specifically dealt with NLP in the political field in the Slovak language, with a specific number of prefixes and suffixes. In this paper, we focus on the area of detection of manipulation of political discussions. The data source consisted of previously collected data in the period from April 2018 with an update until November 19, 2022 realized within this research from the site demagog.sk, which we supplemented with new claims. We trained several classification models on preprocessed data and evaluated them. For multi-class classification, the best results were achieved using logistic regression and a support vector machine trained on the resampled dataset—both achieving an accuracy of 0.56 and a macro F1 score of 0.39. In the case of binary classification, the best results were achieved by logistic regression—accuracy 0.7 and macro F1 score 0.56. These models could help detect manipulation in online political discussions.**

***Index Terms*—Manipulation, fake news detection, automated fact-checking, text analysis**

I. Introduction

Processing non-English texts with standard natural language processing (NLP) tools is largely problematic, mainly from the point of view of their optimization for the English language for which they were mainly designed. The Slovak language is different in the number of prefixes or, on the contrary, suffixes. This situation is well illustrated by the work [1], which is devoted to the analysis of political texts in the Slovak language. For these and other differences in the morphological structure of the analyzed language, it is necessary to modify the general tools [2] in different ways, and only then can their adequate use be possible. Optimization of the tools used is largely dependent on the programmer's experience and knowledge in working with NLP.

For the purposes of the research presented in this paper, we used the dataset by Přibáň et al. [3]. This dataset was created out of the claims of individual politicians in Slovakia, which were then factually verified by expert reviewers associated with the demagog.sk group. The authors do not take responsibility for the actual validity of the opinion of the demagog.sk group, nor for the validity of original claims made by the politicians.

Since the dataset covered findings until April 2018, it was necessary to supplement it with new assessments until November 19, 2022. Using the obtained data, we trained the classification models and compared these models with each other, expecting gradual improvement. We focused on the detection of text characteristics without the use of other, additional data. The work involves processing data and obtaining numerical results.

We address the following research questions:

1) What types of datasets exist and what are their interrelationships in several different stages?
2) How can classification processes be implemented?
3) What are the hardware requirements for neural networks when training models for the needs of different forms and classification techniques?

The key contributions of the paper are as follows:

- Understanding the model creation
- Differences in domain-specific texts
- Overview and understanding of the differences between different models

This rest of the paper is organized as follows. Sections II and III deal with a theoretical summary and understanding of the issue. Sections IV–V cover data collection, dataset construction, model construction, and evaluation. Section VII describes related work. Section VIII concludes the paper.

II. Manipulation of Online Discussion in Antisocial Behavior

Antisocial behavior is generally a type of behavior that is different from what is normally accepted. It includes a whole range of behavior from the practically harmless like coughing in the street ends to the commission of criminal acts. They investigate the manifestation of antisocial behavior by experts from various scientific fields. It is therefore a multidisciplinary issue. A number of definitions associated with antisocial behavior are also based on this fact. In the works, they found the cause in the interplay of internal [4] and external [5] factors acting on the individual during his adolescent development. Manifestations of these factors are transferred to the online environment on social networks.

We looked at the different results within the area to be able to understand how mapping and determining proximity definitions for each of the individual forms can help in the distribution of antisocial behavior (AOB) [6]. This research tries to answer the question of why some forms of antisocial online behavior (AOB) are mistaken for other forms. The team makes frequent mistakes in defining close forms. Thus these mistakes are often impossible to compare the results of the research.

In some countries of the world, there are obligations to removing expressions of AOB from their pages [7]. An example of this law, which is also cited by the mentioned research, is the law from Germany also called NetzDG [8], where companies are obliged to notify at regular intervals (every six months) about the removal of AOB removal. In Germany, there was significant criticism against this law, which demanded its cancellation, or at least its change [9]. But on the other hand, there are positive aspects that justify its existence and help people overcome AOB.

We focus on the area of detection of manipulation of discussions, which constitute a significant part of antisocial behavior in the political field. From the state's point of view and political views on this problem field, it is sometimes also called computational propaganda or information war [10].

Looking at the research area meant taking into account the focus on the detection of different types of AOB in Slovak-speaking groups on social networks [11]–[13]. We also identified the issue of NLP within the text in the Slovak language, which is specific to the number of prefixes and suffixes used in the text. Most tools are optimized for the English language and need to be modified to use them. Optimization is done by practicing Slovak texts, adding stop words, and the like.

III. Natural Language Processing of Slovak

We have already discussed the problem of implementing NLP methods on text in the Slovak language in our earlier work [11], [14]. We pointed out their difference in declension and the number of different grammatical categories.

An extensive survey [15], which used typological information used in the development of various NLP techniques [16], focused on environments in different languages, including Slovak. However, if we consider the number of languages that research has focused on, it is unlikely that even the best programmer will manage to cover all the factors of a language. Specifically, this refers to examples of connecting typological properties, lexical and similar words even for short-text classification in Slovak language [17].

Examples of models [18]–[20] that apply the aforementioned functions to predictions. The FlauBERT [21] model was created on the French language corpus, which was large enough. Another BERT model (Bidirectional Encoder Representations from Transformers), for example, is a model for the Portuguese language created by Brazilian researchers [22] called BERTimbau.

BERT models were also created for other language families. An example is the model for language [23] provides a comprehensive overview of PTMs for NLP. It is generally accepted that it is necessary to take into account the specifics of individual domains. In the beginning, however, it usually starts with models trained on general domains.

However, this assumption is challenged by works such as [24] which argues that for specific domains (such as biomedicine) it is better to start training from scratch on unlabeled texts from the same domain. Research [25] introduces NLP methods for various Indian languages, based on two language families. In the Slovak language, it is important to mention, for example, the SlovakBERT model [26].

IV. Dataset

Our starting point was observation used dataset from Přibáň, Hercig, and Steinberger [3]. The dataset used in this work contains the statements of politicians in the Slovak Republic from the fact-checking website demagog.sk until April 2018. Each claim is assigned to one of the following veracity classes:

1) TRUE
2) FALSE
3) MISLEADING
4) UNVERIFIABLE

The meaning of each class is detailed in [3]. Classes are assigned to claims using fact-checkers from demagog.sk. Each statement also includes the name of the politician who authored it, the party they belong to, and an explanation of why a particular class was chosen for the statement. The dataset contains a total of 12554 records. We did not interfere with the obtained data and all expert results are taken in full. The possibility of a false-positive determination is minimized, as the site evaluates politicians' statements with the help of domain experts [27].

Since we wanted more recent claims as well, we created a crawler that was able to crawl the other 3036 claims (and other information including their classes). We managed to collect data until November 2022.

Finally, we formed the final dataset with which we worked—we took claims and their classes from the original dataset and added the crawled claims and their classes. It contains 15590 records. Table I shows the cardinality of each class. Table II is a preview of the records from the final dataset.

Table I
Cardinality of each class in the final dataset.

TRUE	9857
UNVERIFIABLE	2216
FALSE	2124
MISLEADING	1393

V. Data Preprocessing

We decided to drop records with class UNVERIFIABLE since this class says that it was not possible to verify the veracity of a given claim [27]. In a few experiments, we performed binary classification into classes TRUE and FALSE. In these cases, we also dropped records with class MISLEADING.

Table II
PREVIEW OF THE FINAL DATASET.

Claim SK	Claim EN	Class
Keď sa porovnáme s ostatnými krajinami, Slovenská republika má jednu z najlacnejších cien elektriny aj plynu.	When we compare it with other countries, the Slovak Republic has one of the cheapest prices for both electricity and gas.	TRUE
Teraz Nemci prijali zastropovanie cien plynu až do roku 2024 a títo tu dokážu niečo zagarantovať iba do konca roka.	Now the Germans have accepted a cap on gas prices until 2024, and they can only guarantee something here until the end of the year.	FALSE
Aj najvyšší vládni predstavitelia pán Matovič, pán Kollár a vtedy ešte pán Sulík sa vyjadrili, že to (referendum, pozn.) idú podporiť.	Even the highest government officials, Mr. Matovič, Mr. Kollár, and at that time Mr. Sulík, said that they were going to support it (referendum, note).	MISLEADING
Ruský plyn je momentálne drahší ako je LNG plyn, to povedal aj pán Hirman v relácii.	Russian gas is currently more expensive than LNG gas, Mr. Hirman also said in the show.	UNVERIFIABLE

The following preprocessing steps were applied to each of the remaining claims:

1) punctuation removal,
2) numbers removal,
3) lowercasing,
4) stopwords removal [28],
5) words lemmatization [29]—capitalizes words which are naturally capitalized in the Slovak language (e.g. proper nouns)

As the next step, we transformed preprocessed claims into numerical form. We applied two different types of numerical representation, depending on the classification algorithm the data was put into.

For classification algorithms that were not a neural network, we used TF-IDF vectorization, specifically sklearn's implementation [30].

For recurrent neural networks, we used pretrained word vectors from the following GitHub repository [31]. We used vectors from the file vec-sk-cbow-lemma which contains vectors for lemmas trained using the Continuous Bag of Words algorithm.

VI. Experiments and Results

We trained several classification models on the preprocessed data and evaluated them. We performed multiclass classification to classes TRUE, FALSE, and MISLEADING and binary classification to classes TRUE and FALSE. The dataset suffers from the problem of imbalanced classes (see Table I). We tried to solve this problem by oversampling and also undersampling.

A. Classification algorithms

We used the following classification algorithms (apart from the last algorithm, implementations are from the sklearn library [32]; if not explicitly stated, default sklearn hyperparameters were used):

- Naive Bayes Classifier [33]
- Support Vector Machine [34] (with kernel = 'linear')
- Logistic regression [35] (with max iter=500)
- Random Forrest [36]
- LSTM recurrent neural network

The recurrent neural network was implemented in TensorFlow framework [37]. Its architecture is based on the network proposed by Ivancová [38]. The architecture for multiclass (three-class) classification is shown in Table III. Layer named dense has ReLU activation function and layer named dense_1 has Softmax activation function. The loss function is Categorical cross-entropy and the optimizer is Adam. Architecture for binary classification is the same, except for output layer dense_1—it has only two neurons.

Table III
ARCHITECTURE OF THE LSTM RECURRENT NEURAL NETWORK FOR THREE-CLASS CLASSIFICATION.

Layer (type)	Output Shape	Param #
embedding (Embedding)	(None, None, 200)	3000400
lstm (LSTM)	(None, 128)	168448
dense (Dense)	(None, 128)	16512
dense_1 (Dense)	(None, 3)	387
Total params: 3,185,747		
Trainable params: 185,347		
Non-trainable params: 3,000,400		

All models were trained on the following hardware: CPU Intel Core i5-9300HF, RAM 16 GB and GPU NVIDIA GeForce GTX 1650 4GB. The GPU was used in the case of the neural network and the CPU in the case of other models.

B. Experiments

For each algorithm, we trained a multiclass classifier and a binary classifier on an imbalanced dataset, oversampled dataset, and undersampled dataset. Metrics used for performance evaluation of the models are accuracy and macro F1 score. We performed 10-fold cross-validation for models trained on the imbalanced and oversampled dataset (except for the recurrent neural network). In other cases, we just split the data into training and validation sets in a ratio of 80:20 and used these sets for training and validation.

C. Results

We provide the results of our experiments in the following tables (results for the neural network are in a separate table):

- results for 10-fold cross-validation on an imbalanced dataset in Table IV
- results for 10-fold cross-validation on an oversampled dataset in Table V
- results for validation on an undersampled dataset in Table VI
- results for the recurrent neural network in Table VII

As we can see in the tables, in the case of multiclass classification the best results were achieved by logistic regression and

Table IV
CLASSIFICATION ON IMBALANCED DATASET—10-FOLD CROSS-VALIDATION.

Classification algorithm	Accuracy	Macro F1
Naive Bayes - multiclass classification	0.74	0.28
Naive Bayes - binary classification	0.82	0.45
SVM - multiclass classification	0.74	0.29
SVM - binary classification	0.82	0.46
Logistic regression - multiclass classification	0.74	0.3
Logistic regression - binary classification	0.82	0.46
Random Forest - multiclass classification	0.74	0.32
Random Forest - binary classification	0.82	0.49

Table V
CLASSIFICATION ON OVERSAMPLED DATASET—10-FOLD CROSS-VALIDATION.

Classification algorithm	Accuracy	Macro F1
Naive Bayes - multiclass classification	0.5	0.38
Naive Bayes - binary classification	0.67	0.55
SVM - multiclass classification	0.56	0.39
SVM - binary classification	0.69	0.56
Logistic regression - multiclass classification	0.56	0.39
Logistic regression - binary classification	0.7	0.56
Random Forest - multiclass classification	0.74	0.32
Random Forest - binary classification	0.82	0.51

Table VI
CLASSIFICATION ON THE UNDERSAMPLED DATASET—NO CROSS-VALIDATION.

Classification algorithm	Accuracy	Macro F1
Naive Bayes - multiclass classification	0.35	0.31
Naive Bayes - binary classification	0.58	0.51
SVM - multiclass classification	0.4	0.33
SVM - binary classification	0.56	0.5
Logistic regression - multiclass classification	0.39	0.32
Logistic regression - binary classification	0.57	0.51
Random Forest - multiclass classification	0.39	0.33
Random Forest - binary classification	0.54	0.49

Table VII
CLASSIFICATION WITH LSTM NEURAL NETWORK—NO CROSS-VALIDATION; RESULTS AFTER 10 EPOCHS.

LSTM neural network	Accuracy	Macro F1
Imbalanced dataset - multiclass classification	0.74	0.28
Imbalanced dataset - binary classification	0.82	0.45
Oversampled dataset - multiclass classification	0.1	0.06
Oversampled dataset - binary classification	0.18	0.15

support vector machine trained on oversampled dataset—they both achieved accuracy 0.56 and macro F1 score 0.39.

In the case of binary classification, the best results were achieved by logistic regression trained on oversampled dataset—accuracy 0.7 and macro F1 score 0.56.

Our results are comparable with results achieved in different works which dealt with classification of short-text claims.

When we compare our work with Přibáň, Hercig, and Steinberger's [3] work (the work in which the authors provide the original dataset and perform similar experiments to ours), we tried different text preprocessing and also tried multiple different classification algorithms. Our classification to three classes achieved better accuracy by 0.21 and better macro F1 score by 0.04 in comparison to their four class classification. In the case of binary classification, we managed to improve accuracy by 0.12 but the macro F1 score dropped by 0.02.

These models can serve as a tool for detecting manipulations in political discussions, e.g., for the classification of the texts written by Slovak politicians.

VII. RELATED WORK

In the practical part of our work, we deal with detection of false information in short-text claims of Slovak politicians. This section explores similar solutions.

Wang [39] presented the dataset LIAR which consists of 12836 English fact-checked short-text claims from the fact-checking website politifact.com. Each claim has various meta-data associated with it (speaker, context, subject…) and is assigned 1 of the 6 following classes: true, mostly-true, half-true, barely true, false, and pants-fire. Wang performed various experiments where he tried to predict the class of each claim using traditional machine learning algorithms and also LSTM and convolutional neural networks. He conducted experiments where only textual features were taken into account and also experiments where the metadata were considered. The best accuracy achieved on the validation set was achieved with a convolutional neural network which took text of the claim and speaker (as metadata) into account—27.7%. The best accuracy achieved on the test set was achieved again with a convolutional neural network which took text and all available metadata into account—27.4%.

Researchers tried to improve the classification accuracy on the LIAR dataset with LSTM and CNN based models and also by adding new metadata to the claims [40]. Kirillin and Strube achieved accuracy of 45.7% by including a credibility vector of the speaker. This vector contains counts of speaker's claims from each class [40], [41].

Přibáň, Hercig and Steinberger [3] performed similar experiments to ours. We also used the same dataset (and added new data to it). The dataset contains fact-checked short-text claims of Slovak politicians from fact-checking website demagog.sk. The authors performed multiclass classification to classess true, false, misleading and unverifiable and also binary classification to classess true and false. In the case of multiclass classification to all four classes both accuracy and macro F1 score were 0.35. For binary classification, the researchers achieved value 0.58 for both accuracy and macro F1 score. They also performed experiments where they added text explaining why the particular class was assigned to the given claim to the text of the claim. These explanations are part of the dataset. Researchers were able to achieve better results in these cases.In the case of multiclass classification both accuracy and macro F1 score were increased by 0.43. In the case of binary classification both accuracy and macro F1 score were increased by 0.27.

Ivancová, Sarnovský and Maslej-Krešňáková [38] tried to detect Slovak fake news articles with convolutional neural network and LSTM neural network. CNN achieved accuracy of 92,38% and LSTM 93,56%.

VIII. Conclusions and Future Work

The aim of the research reported in this paper was to investigate the classification of AOB in Slovak-speaking groups on social networks. We focused on detecting the manipulation of discussions, which constitutes a significant part of antisocial behavior in the political field. Through training the model, we were able to better understand and train several classification models on the preprocessed data and evaluated them. In the case of multiclass classification, the best results were achieved by logistic regression and support vector machine trained on an oversampled dataset.

Our future research will concern the implementation of the results to detect the manipulation of discussions in the political domain on social networks.

Acknowledgments

We want to thank Pavle Dakić for his support in writing this paper by providing insights and suggestions for improving the obtained results.

The work reported here was supported by the Slovak national project Increasing Slovakia's Resilience Against Hybrid Threats by Strengthening Public Administration Capacities (Zvýšenie odolnosti Slovenska voči hybridným hrozbám pomocou posilnenia kapacít verejnej správy) (ITMS code: 314011CDW7), co-funded by the European Regional Development Fund (ERDF), the Operational Programme Integrated Infrastructure for the project: Research in the SANET network and possibilities of its further use and development (ITMS code: 313011W988), co-funded by the ERDF, and by the Slovak Research and Development Agency under the contract No. APVV-15-0508.

References

[1] R. Štefančík and E. Stradiotová, "The concept of nation in the language of the Slovak right-wing extremists," *Journal of Comparative Politics*, vol. 14, no. 2, jul 2021.

[2] N. Hroncova and P. Dakic, "Research study on the use of CI/CD among slovak students," in *2022 12th International Conference on Advanced Computer Information Technologies (ACIT)*. IEEE, sep 2022.

[3] P. Přibáň, T. Hercig, and J. Steinberger, "Machine learning approach to fact-checking in west slavic languages," in *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, 2019, pp. 973–979.

[4] N. J. Kolla and C. C. Wang, "Alcohol and violence in psychopathy and antisocial personality disorder: Neural mechanisms," in *Neuroscience of Alcohol*. Elsevier, 2019, pp. 277–285.

[5] L. O'Malley and S. Grace, "Social capital and co-location: A case study of policing anti-social behaviour," *International Journal of Police Science and Management*, vol. 23, pp. 306–316, may 2021.

[6] A. Hrčková, I. Srba, R. Móro, R. Blaho, J. Šimko, P. Návrat, and M. Bieliková, "Unravelling the basic concepts and intents of misbehavior in post-truth society," *Biblioteca Anales de Investigación*, vol. 15, no. 3, pp. 421–428, jan 2019.

[7] E. Zinovyeva, W. K. Härdle, and S. Lessmann, "Antisocial online behavior detection using deep learning," *Decision Support Systems*, vol. 138, pp. –, jun 2020.

[8] R. js/amp (AFP, AP, "Germany fines facebook over hate speech complaints," https://www.dw.com/en/germany-fines-facebook-for-underreporting-hate-speech-complaints/a-49447820, 2019, accessed: 2022-01-21.

[9] A. Pearson, "Fdp, greens and left call to replace hate speech law," https://www.dw.com/en/german-opposition-parties-call-to-replace-online-hate-speech-law/a-42058030, 2018, accessed: 2022-01-21.

[10] E. Ferrara, H. Chang, E. Chen, G. Muric, and J. Patel, "Characterizing social media manipulation in the 2020 u.s. presidential election," *First Monday*, vol. 25, no. 11, oct 2020.

[11] I. Stupavský and V. Vranić, "A study of media texts in the Slovak language," in *2022 IEEE Zooming Innovation in Consumer Technologies Conference (ZINC)*. IEEE, may 2022.

[12] M. Krocka, P. Dakic, and V. Vranic, "Extending parking occupancy detection model for night lighting and snowy weather conditions," in *2022 IEEE Zooming Innovation in Consumer Technologies Conference (ZINC)*. IEEE, may 2022.

[13] T. Golis, P. Dakić, and V. Vranić, "Creating microservices and using infrastructure as code within the CI/CD for dynamic container creation," in *2022 IEEE 16th International Scientific Conference on Informatics (Informatics)*. IEEE, nov 2022.

[14] I. Stupavský and V. Vranić, "Analysing the controversial social media community," in *2022 IEEE 16th International Scientific Conference on Informatics (Informatics)*, 2022, pp. 299–303.

[15] E. M. Ponti, H. O'Horan, Y. Berzak, I. Vulić, R. Reichart, T. Poibeau, E. Shutova, and A. Korhonen, "Modeling language variation and universals: A survey on typological linguistics for natural language processing," 2018.

[16] X. Chen, A. H. Awadallah, H. Hassan, W. Wang, and C. Cardie, "Multi-source cross-lingual model transfer: Learning what to share," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 2019.

[17] M. Borecky and J. Lang, "Ontology-driven detection of redundancy in short texts and its visualization," in *Proceedings of the 23rd International Conference on Computer Systems and Technologies*, ser. CompSysTech '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 103–107. [Online]. Available: https://doi.org/10.1145/3546118.3546149

[18] M. Kročka, P. Dakić, and V. Vranić, "Automatic license plate recognition using OpenCV," in *2022 12th International Conference on Advanced Computer Information Technologies (ACIT)*. IEEE, sep 2022.

[19] R. Szarka, P. Dakic, and V. Vranic, "Cost-effective real-time parking space occupancy detection system," in *2022 12th International Conference on Advanced Computer Information Technologies (ACIT)*. IEEE, sep 2022.

[20] P. Dakić and M. Živković, “An overview of the challenges for developing software within the field of autonomous vehicles,” in *7th Conference on the Engineering of Computer Based Systems*. ACM, may 2021.
[21] H. Le, L. Vial, J. Frej, V. Segonne, M. Coavoux, B. Lecouteux, A. Allauzen, B. Crabbé, L. Besacier, and D. Schwab, “Flaubert: Unsupervised language model pre-training for french,” 2019.
[22] F. Souza, R. Nogueira, and R. Lotufo, “BERTimbau: Pretrained BERT models for brazilian portuguese,” in *Intelligent Systems*. Springer International Publishing, 2020, pp. 403–417.
[23] X. Qiu, T. Sun, Y. Xu, Y. Shao, N. Dai, and X. Huang, “Pre-trained models for natural language processing: A survey,” 2020.
[24] Y. Gu, R. Tinn, H. Cheng, M. Lucas, N. Usuyama, X. Liu, T. Naumann, J. Gao, and H. Poon, “Domain-specific language model pretraining for biomedical natural language processing,” 2020.
[25] D. Kakwani, A. Kunchukuttan, S. Golla, G. N.C., A. Bhattacharyya, M. M. Khapra, and P. Kumar, “IndicNLPSuite: Monolingual corpora, evaluation benchmarks and pre-trained multilingual language models for indian languages,” in *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, 2020.
[26] M. Pikuliak, Š. Grivalský, M. Konôpka, M. Blšták, M. Tamajka, V. Bachratý, M. Šimko, P. Balážik, M. Trnka, and F. Uhlárik, “SlovakBERT: Slovak masked language model,” 2021.
[27] Demagog.sk, “Ako pracujeme,” https://demagog.sk/ako-pracujeme, accessed: 2023-01-20.
[28] M. Prata, “Slovak stop words w2v,” https://www.kaggle.com/code/mpwolke/slovak-stop-words-w2v/data?select=slovak.txt, 2021, accessed: 2023-01-20.
[29] S. N. Corpus, “Morphologic database of slovak language,” https://korpus.sk/en/corpora-and-databases/databases/morphology-database/, 2015, accessed: 2023-01-20.
[30] Scikit-learn.org, “sklearn.feature_extraction.text.tfidfvectorizer,” https://scikit-learn.org/stable/modules/generated/sklearn.feature_extraction.text.TfidfVectorizer.html, 2022, accessed: 2023-01-20.
[31] B. Juraj, “word2vec-sk,” https://github.com/essential-data/word2vec-sk, 2015, accessed: 2023-01-20.
[32] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
[33] Scikit-learn.org, “sklearn.naive_bayes.multinomialnb,” https://scikit-learn.org/stable/modules/generated/sklearn.naive_bayes.MultinomialNB.html#sklearn.naive_bayes.MultinomialNB, 2022, accessed: 2023-01-20.
[34] ——, “sklearn.svm.svc,” https://scikit-learn.org/stable/modules/generated/sklearn.svm.SVC.html, 2022, accessed: 2023-01-20.
[35] ——, “sklearn.linear_model.logisticregression,” https://scikit-learn.org/stable/modules/generated/sklearn.linear_model.LogisticRegression.html#sklearn.linear_model.LogisticRegression, 2022, accessed: 2023-01-20.
[36] ——, “sklearn.ensemble.randomforestclassifier,” https://scikit-learn.org/stable/modules/generated/sklearn.ensemble.RandomForestClassifier.html, 2022, accessed: 2023-01-20.
[37] M. Abadi *et al.*, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015, software available from tensorflow.org. [Online]. Available: https://www.tensorflow.org/
[38] K. Ivancová, M. Sarnovský, and V. Maslej-Krešňáková, “Fake news detection in slovak language using deep learning techniques,” in *2021 IEEE 19th World Symposium on Applied Machine Intelligence and Informatics (SAMI)*. IEEE, 2021, pp. 000 255–000 260.
[39] W. Y. Wang, "“Liar, Liar Pants on Fire”: A New Benchmark Dataset for Fake News Detection,” *arXiv preprint arXiv:1705.00648*, 2017.
[40] R. Oshikawa, J. Qian, and W. Y. Wang, “A survey on natural language processing for fake news detection,” *arXiv preprint arXiv:1811.00770*, 2018.
[41] A. Kirilin and M. Strube, “Exploiting a speaker’s credibility to detect fake news,” in *Proceedings of Data Science, Journalism & Media workshop at KDD (DSJM’18)*, 2018.