

Testing Pre-trained Transformer Models for Lithuanian News Clustering

Lukas Stankevičius
Faculty of Informatics
Kaunas University of Technology
Kaunas, Lithuania
lukas.stankevicius@ktu.edu

Mantas Lukoševičius
Faculty of Informatics
Kaunas University of Technology
Kaunas, Lithuania
mantas.lukosevicius@ktu.edu

Abstract—A recent introduction of Transformer deep learning architecture made breakthroughs in various natural language processing tasks. However, non-English languages could not leverage such new opportunities with the English text pre-trained models. This changed with research focusing on multilingual models, where less-spoken languages are the main beneficiaries. We compare pre-trained multilingual BERT, XLM-R, and older learned text representation methods as encodings for the task of Lithuanian news clustering. Our results indicate that publicly available pre-trained multilingual Transformer models can be fine-tuned to surpass word vectors but still score much lower than specially trained doc2vec embeddings.

Index Terms—document clustering, document embedding, Lithuanian news articles, Transformer model, BERT, XLM-R, multilingual

I. INTRODUCTION

Appearance of a novel Transformer deep learning architecture [1] sparked a rapid research progress in the Natural Language Processing (NLP) field. Table I clearly depicts how quickly models reached human performance on popular NLP evaluation datasets. In less than two years after publication, two evaluation datasets, SQuAD2.0 [2] and GLUE [3] had human performance outmatched. Currently every top scoring model is of Transformer architecture. The situation was not changed by a newer SuperGLUE [4] task set which had been left with only a tiny gap to human performance. These datasets were among the most popular to evaluate new Transformer models and showed the effectiveness of this new architecture.

TABLE I
DIFFICULTY OF THE MOST POPULAR NLP EVALUATION DATASETS

Dataset		Score			
Name	Year	Initial	Current	Human	Type
RACE [5]	2017	44.1	89.4	94.5	Accuracy
SQuAD2.0 [2]	2018	66.3	92.58	89.542	F1
GLUE [3]	2019	70.0	90.3	87.1	Average
SuperGLUE [4]	2019	71.5	89.3	89.8	Average

There is a need to create NLP models for less-spoken languages. Apart from being less popular than English or Chinese, less-spoken languages also have less content to train the models on. Just the top-10 out of 6000 languages in use

today make up 76.3% of the total content on the internet¹. Such situation encourages not only to pursue creation of NLP models for other languages but also to look for ways to transfer the knowledge from the content-rich language models.

The most common way to satisfy this need for less spoken languages is to pre-train multilingual models. Examples are LASER [6] (93 languages), multilingual BERT [7] (104 languages) and XLM-R [8] (100 languages). The authors of XLM [9] showed that training Nepali language model on Wikipedia together with additional data from both English and Hindi decreased perplexity to 109.3 on Nepali, compared to single Nepali training perplexity of 157. Transfer learning and zero-shot translation between language pairs never seen explicitly during training was shown to be possible in [10]. Overall, multilingual models can cover many languages, be trained without any cross-lingual supervision and use the bigger languages to benefit the smaller ones.

Lithuanian language does not yet have BERT-scale monolingual NLP model. It is relatively very little spoken in a world. However, as a national language of one of European Union member states, Lithuanian is usually included in the most of pre-trained multilingual models. The aim of this work is to use such Transformer-type models to generate text embeddings and evaluate them on clustering of Lithuanian news articles. Specifically, we will use well known baselines – multilingual BERT and recently published XLM-R, trained on more than two terabytes of filtered CommonCrawl data.

We chose clustering task to also try to advance the field of data mining. The surge of information, particularly news data, demands tools to help users to “analyze and digest information and facilitate decision making” [11]. Unlike classification, clustering is universal in that it can handle unknown categories. Therefore it is well suited for the quickly changing news articles data.

II. LITERATURE REVIEW

In this section we review the two first consecutive phases of common natural language processing (NLP) tasks: text preprocessing and text representation [11]. These stages recently are of the most active research and culminated in the development

¹<https://www.internetworldstats.com/stats7.htm>

of the Transformer architecture. We also examine relevant NLP contributions for Lithuanian language and our task of news clustering.

A. Text Preprocessing

Text preprocessing involves selection of features that will bear the understanding of text. The most elementary approach is tokenization into simple characters or words. The finer the tokenization, the smaller is the resulting vocabulary and the more challenging task is given to the NLP model. On the other hand, coarser tokens drastically increase vocabulary size and induce other problems such as sparseness. The middle ground is statistically significant n-grams of both words and chars. Examples of this type of tokenizers are SentencePiece [12], BPE [13], and WordPiece [14]. They are often used in the state-of-the-art (SOTA) Transformer models and are shipped together with publicly available pre-trained models. This way manual tokenization step is skipped.

There are number of methods to filter word level tokens. It includes lowercasing, stemming, lemmatization, filtering by maximum and minimum document frequencies (ignoring tokens that are too rare or too common throughout the documents). However, it was shown in [15] that such filtering benefits only the classical text representation approaches such as *tf-idf*, while shallow neural network models, doc2vec [16], benefited of not using any such filtering.

B. Text Representation

Although tokenized text remains meaningful to us, models still can not operate on it directly. They need it in a numerical form. The preferable way is to derive vector representation for each text sample. Cosine similarity is the simplest example of models operating on (comparing) these embeddings. The classical approach to text representation uses a Bag of Words (BoW) model. As the name suggests, the order of tokens is lost here and each document is represented by bare counts (histogram) of its tokens. Therefore token weighting such as *tf-idf* is involved. The higher weight of *tf-idf* is, the more descriptive token for a given document is. Given the number of word w occurrences in a document d as $tf_{w,d}$, number of documents containing word w as df_w and total number of documents N , $tf-idf_{w,d}$ is given by

$$tf-idf_{w,d} = tf_{w,d} \cdot \log(N/df_w). \quad (1)$$

BoW approaches suffer from several problems. The vector length for each document is the same as the size of the vocabulary. Typically, the vocabulary size is huge and this induces major memory constraints. The embedded vectors are also very sparse as each document uses only the small subset of the vocabulary. Various methods, such as Latent Semantic Analysis (LSA) using Singular Value Decomposition (SVD), are employed to reduce the dimensionality. Nevertheless, SVD has to operate on the same high dimensional *documents* \times *tokens* matrix.

Work of [17] revolutionised word embedding calculations. Previous word embeddings, known as co-occurrence vectors,

were superseded. They were calculated as direct probabilities of surrounding words in a context window of a given length. The new word2vec [17] algorithm uses the same training inputs, except the goal is not to calculate the word distribution, but to derive such embedding weights that context words would be predicted with maximum accuracy. Such setup significantly reduced the word vector size and eliminated problems of high dimensionality and sparseness. Later, the next word embedding model, Global Vectors (GloVe) for word representation was presented [18]. It merged advantages of both the matrix factorization and the shallow window-based methods. Currently it is the most used method for independent word vectors.

Original word2vec word vectors were extended to paragraph vectors by doc2vec model [16]. Here each sequence of tokens has its own embedding in the same space as that of words. CBOW and Skip-gram architectures from the original word2vec algorithm were applied to documents correspondingly as PV-DM and PV-DBOW. Distributed Memory model of Paragraph Vector (PV-DM) is tasked to predict the next context word given the previous context and document vector, thus the vector has to sustain memory of that is missing. Distributed Bag of Words version of Paragraph Vectors (PV-DBOW) is forced to predict context vectors randomly, given only the document vector. The original work [16] used a concatenation of the both models. However, in [15] it was found that PV-DBOW alone gave the best results.

The most recent text representation models produce contextualised token vectors. Models like ELMO [19] and various Transformers have each of their inputs to interact with the other ones. This leads to each token vector being aware of the others. Contextualisation solved the problem of word polysemy seen in word2vec or GloVe models.

The Transformer architecture presented in 2017 [1] excels other contextualised models due to several reasons. Firstly, only after a single layer each input representation becomes aware of the other ones. For Recurrent Neural Networks (RNN) like ELMO it took n layers, where n is the sequence length. Another advantage over recurrent architectures is that Transformer is very parallelizable. It does not need to wait for a hidden state of the previous word as is the case with RNN. This particular feature led to creation of multi-billion-parameter Transformers such as GPT-2 [20], T5 [21], Megatron [22], and T-NLG². Despite huge success of Transformer models, it can not process long sequences as its complexity per layer is $O(n^2 \cdot d)$ where d is representation dimension. For example, the maximum input length of the popular BERT model is just 510 tokens. This and other problems of Transformer architecture currently are researched very actively.

C. Related Work on Lithuanian Language

There are several works on Lithuanian text clustering. [23] used internet daily newspaper articles from the *Lrytas.lt* news

²<https://www.microsoft.com/en-us/research/blog/turing-nlg-a-17-billion-parameter-language-model-by-microsoft/>

website and information from the largest internet forum for mothers – *supermama.lt*. k -means and Expectation Maximization (EM) algorithms were compared on BoW data representation. It was found that optimal clustering algorithm is k -means with cosine similarity. Other work analysed unsupervised feature selection for document clustering [24]. Authors found that $tf-idf$ weighting with 3000 features and spherical k -means clustering algorithm works best. A similar observation is expressed in [25]. Here k -means was compared to various hierarchical clustering algorithms and was not superseded. Authors found that $tf-idf$ together with stemming is superior to other approaches.

[15] compared BoW and doc2vec (PV-DBOW) Lithuanian news article representations for document clustering. It was shown that PV-DBOW representation, trained on the whole dataset is superior to the BoW method. Authors also investigated various PV-DBOW hyperparameters and outlined recommendations for training method weights.

There is also other NLP work on Lithuanian language. [26] and [27] compared traditional and deep learning approaches for Lithuanian internet comments sentiment classification. Authors demonstrated that traditional Naïve Bayes Multinomial and Support Vector Machine methods outperformed LSTM and Convolutional neural networks. Other work [28] compared CBOw and Skip-gram word embedding architectures and found the first one to be superior. We can add that in this current work we noticed a similar tendency for document vectors: equivalent version of PV-DBOW in our initial experiments outperformed PV-DMM architecture.

We have not found any previous work on Lithuanian language using Transformer models.

III. THE DATA

We followed methodology of [15] and expanded their dataset from 82 793 up to 260 146 articles. Although average number of characters in each our text sample is 2 948, several scraped articles were very small and resulted in empty vectors during averaging the GloVe vectors. Due to this reason we filtered all articles with less than 200 characters and this resulted in a final dataset of 259 996 texts.

The data consist of Lithuanian news articles scraped from *lrt.lt*, *15min.lt*, and *delfi.lt* websites. The number of texts are correspondingly 26 344, 133 587, and 100 065. Due to the absence of sitemap in *lrt.lt* website, we did not scrape more articles from this site than is already scraped in [15].

Evaluation of clustering requires existing knowledge of the potential clusters. For this task we leveraged article category labels extracted from each article URL. Following the mappings of [25], the labels were unified from over a hundred categorical descriptions down to 12 distinct categories. The resulting categories of the articles are:

- Lithuanian news (60 158 articles);
- World news (68 635 articles);
- Crime (30 967 articles);
- Business (19 964 articles);
- Cars (6 313 articles);

- Sports (14 910 articles);
- Technologies (4 438 articles);
- Opinions (9 728 articles);
- Entertainment (2 462 articles);
- Life (3 811 articles);
- Culture (7 967 articles);
- Other (30 643 articles that do not fall into the previous categories).

During most of experiments we employed a smaller subset of the dataset described above. We sampled randomly 125 news articles from each of the 12 categories. That results in a total of 1 500 articles equally distributed among the categories and corresponds to the data required for one clustering. We planned to make 50 independent clusterings to average results and enhance their reproducibility. Thus we repeated independent sampling of 1 500 equally distributed articles and used a subset of 55 487 unique articles (with repetitions it would be 75 000). During each experiment we calculated embedding vectors only for those 55 487 news articles.

IV. METHODS

A. Clustering

We use k -means clustering algorithm. Due to its high speed, it is suitable for large corpora [23] and outperforms other clustering algorithms [25]. During experiments we feed vectorized document representations and the expected number of clusters k to k -means and receive document assignments to clusters. We set each of 50 k -means initialisations the same.

B. Evaluation

First, we calculate the following confusion matrix elements:

- TP – pairs of articles that have same category label and are predicted to be in the same cluster;
- TN – pairs of articles that belong to different categories and are predicted to be in different clusters;
- FP – pairs of articles that belong to different categories but are predicted to be in the same cluster;
- FN – pairs of articles that have same category label but are predicted to be in different clusters.

We chose to evaluate clusters by Matthews Correlation Coefficient (MCC) score due to its reliability as described in [29]. It is calculated as

$$MCC = \frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}. \quad (2)$$

The MCC score ranges from -1 to 1. Scores around 0 value correspond to random clustering, while close to 1 indicate perfect matching.

C. PV-DBOW

This doc2vec version was trained on all our dataset – total of 259 996 Lithuanian news articles. We preprocessed the dataset by lowercasing and tokenizing it into words. The same vector size (100), number of epochs (10), window (12), and minimum count (4) parameters were used as in [15]. PV-DBOW returns a single embedding for each document so no further aggregation is required.

D. GloVe

We performed our own text preprocessing during experiments with GloVe [18] type Lithuanian word vectors [30]. It combined lowercasing and word level tokenization. Out of whole unique 1 028 816 tokens from our whole dataset 311 470 were also present in Lithuanian GloVe vectors. This unique tokens intersection amounts to 30% of our tokens and up to 94% of GloVe.

We tried several ways of aggregating GloVe vectors:

- calculating an average of all the word vectors in the article;
- weighting all tokens with *tf-idf* and calculating an average of the 20 word vectors with the highest weight;
- weighting all tokens in an article with *Softmax(tf-idf)* and calculating a weighted average of all the word vectors in the article.

E. Multilingual BERT

BERT [7] outputs the same number of vectors as it is fed inputs. The first is a special [CLS] token which is designed to be used in sentence level tasks. The following are text data tokens, ending with the last [SEP] token. Optionally, one can add a second [SEP] token in a middle of an input sequence to separate two text segments. In our experiments we input only one segment and always separately try the [CLS] token output vector and the averages of all token vectors.

Pre-trained models like multilingual BERT are supposed to be fine-tuned for the desirable task. We performed Masked Language Modelling fine-tuning on half of our subset data – 27 743 news articles. We trained with batch size of 4 for 5 epochs totalling 68 505 steps for uncased and 75 985 steps for cased versions of pre-trained multilingual BERT model.

The maximum number of input tokens to BERT is 512, including special tokens. The most articles are within 512 tokens limit but some are longer. We tried to estimate effect of this constraint by trying to feed even fewer tokens and analyzing how mean MCC score changes with the longer input sequence.

We carried out our experiments in Google Colab³ environment. It offers 12 GB of RAM and GPU-accelerated machines which allows an order of magnitude speed up of BERT model compared to CPU.

F. XLM-R

XLM-R [8] is one of the recent multilingual language models, much bigger than multilingual BERT. It is trained on 2 terabytes of filtered text from which 13.7 GB is Lithuanian. The huge size of this model limited our experiments. We only calculated outputs of the first 512 tokens for each news article. It took approximately a total of 40 hours.

³<https://colab.research.google.com/>

V. RESULTS

A. GloVe

Results with aggregation of GloVe vectors are presented in Table II. It is clearly seen that applying *tf-idf* weighting to select the best tokens to average can significantly surpass the simple average of all vectors.

TABLE II
METHODS OF COMBINING GLOVE VECTORS

Vector aggregation	MCC score	
	Mean	Std
Average	0.203	0.016
Softmax(tf-idf) weighted average	0.264	0.017
Average of 20 highest tf-idf tokens	0.264	0.024

B. Multilingual BERT

The effect of multilingual BERT fine-tuning on Lithuanian news articles is depicted in Fig. 1. One can clearly see that (1) the fine-tuning improves the clustering results, (2) the average of all tokens is much better than only the [CLS] vector, and (3) the uncased model version outperforms the cased one.

To our surprise, the best results were obtained with limiting the number of tokens to only the first 144 (see Fig. 2). This can be attributed to the more important information being in the beginning of news article. We observed the same tendency with the XLM-R model.

C. The Best Models

We tried four different models to represent Lithuanian news articles. Do Transformer models score better than PV-DBOW? As can be seen in Table III, the best Transformer model managed to outperform GloVe vectors. However, PV-DBOW model is far ahead with the mean MCC score of 0.442.

TABLE III
COMPARISON OF METHODS

Text representation method	MCC score	
	Mean	Std
PV-DBOW	0.442	0.028
Uncased fine-tuned BERT, average of first 144 tokens	0.322	0.020
Softmax(tf-idf) weighted average of GloVe vectors	0.264	0.017
XLM-R, average of first 288 tokens, total fed 512	0.251	0.016

VI. CONCLUSIONS

In this work we compared multilingual BERT, XLM-R, GloVe, and PV-DBOW text representations for Lithuanian news clustering. For BERT we found out that the average of only the first 144 token vectors outperforms longer aggregations or the [CLS] token vector. We observed that BERT fine-tuning with Lithuanian news articles also improves the results. The other pre-trained Transformer type model XLM-R was too computationally expensive to optimize and out of the four its initial representations scored the worst.




Fig. 1. MCC score dependence on multilingual BERT model type (cased or uncased), language modelling fine-tuning steps, and token vector aggregation method (average of all first 128 tokens or just the first [CLS] token). Markers show mean of MCC and shadows – standard deviation.




Fig. 2. MCC score dependence on the number of first tokens used and the token vector aggregation method (average of all selected tokens or just the first [CLS] token) for multilingual uncased BERT model. Markers show mean of MCC and shadows – standard deviation.

Regarding GloVe vectors, we found that its best Softmax(tf-idf) embeddings (mean MCC score 0.264) are outperformed by the BERT. Nevertheless, the best text representation method proved to be PV-DBOW with mean MCC score 0.442. Our work on generating representations for Lithuanian news clustering showed that multilingual pre-trained Transformers can be better than independent GloVe vectors but under-performs against specially trained simpler PV-DBOW.

Multilingual BERT MCC score kept rising till last fine-

tuning steps and it is not clear how large the improvement could be accomplished training longer. Our future plan is to clarify this by using more data. We also plan to train a new monolingual BERT model specifically for Lithuanian language. It would be interesting to know if these resource-“hungry” approaches could surpass the score of the relatively simple PV-DBOW method.

REFERENCES

- [1] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [2] P. Rajpurkar, R. Jia, and P. Liang, “Know what you don’t know: Unanswerable questions for squad,” *arXiv preprint arXiv:1806.03822*, 2018.
- [3] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman, “Glue: A multi-task benchmark and analysis platform for natural language understanding,” *arXiv preprint arXiv:1804.07461*, 2018.
- [4] A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, O. Levy, and S. Bowman, “Superglue: A stickier benchmark for general-purpose language understanding systems,” in *Advances in Neural Information Processing Systems*, 2019, pp. 3261–3275.
- [5] G. Lai, Q. Xie, H. Liu, Y. Yang, and E. Hovy, “Race: Large-scale reading comprehension dataset from examinations,” *arXiv preprint arXiv:1704.04683*, 2017.
- [6] M. Artetxe and H. Schwenk, “Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond,” *Transactions of the Association for Computational Linguistics*, vol. 7, pp. 597–610, 2019.
- [7] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [8] A. Conneau, K. Khandelwal, N. Goyal, V. Chaudhary, G. Wenzek, F. Guzmán, E. Grave, M. Ott, L. Zettlemoyer, and V. Stoyanov, “Unsupervised cross-lingual representation learning at scale,” *arXiv preprint arXiv:1911.02116*, 2019.
- [9] A. Conneau and G. Lample, “Cross-lingual language model pretraining,” in *Advances in Neural Information Processing Systems*, 2019, pp. 7057–7067.

- [10] M. Johnson, M. Schuster, Q. V. Le, M. Krikun, Y. Wu, Z. Chen, N. Thorat, F. Viégas, M. Wattenberg, G. Corrado *et al.*, “Google’s multilingual neural machine translation system: Enabling zero-shot translation,” *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 339–351, 2017.
- [11] C. C. Aggarwal and C. Zhai, *Mining text data*. Springer Science & Business Media, 2012.
- [12] T. Kudo and J. Richardson, “Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing,” *arXiv preprint arXiv:1808.06226*, 2018.
- [13] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” *arXiv preprint arXiv:1508.07909*, 2015.
- [14] Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey *et al.*, “Google’s neural machine translation system: Bridging the gap between human and machine translation,” *arXiv preprint arXiv:1609.08144*, 2016.
- [15] L. Stankevičius, “Clustering of Lithuanian news articles using document embeddings,” Master’s thesis, Kaunas University of Technology, 2019.
- [16] Q. Le and T. Mikolov, “Distributed representations of sentences and documents,” in *International conference on machine learning*, 2014, pp. 1188–1196.
- [17] T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” *arXiv preprint arXiv:1301.3781*, 2013.
- [18] J. Pennington, R. Socher, and C. D. Manning, “Glove: Global vectors for word representation,” in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.
- [19] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer, “Deep contextualized word representations,” *arXiv preprint arXiv:1802.05365*, 2018.
- [20] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever, “Language models are unsupervised multitask learners,” *OpenAI Blog*, vol. 1, no. 8, p. 9, 2019.
- [21] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, “Exploring the limits of transfer learning with a unified text-to-text transformer,” *arXiv preprint arXiv:1910.10683*, 2019.
- [22] M. Shoenberger, M. Patwary, R. Puri, P. LeGresley, J. Casper, and B. Catanzaro, “Megatron-lm: Training multi-billion parameter language models using gpu model parallelism,” *arXiv preprint arXiv:1909.08053*, 2019.
- [23] G. Ciganaitė, A. Mackutė-Varoneckienė, and T. Krilavičius, “Text documents clustering.” *Informacinės technologijos : 19-oji tarpuniversitetinė tarptautinė magistrantų ir doktorantų konferencija "Informacinė visuomenė ir universitetinės studijos" (IVUS 2014) : konferencijos pranešimų medžiaga*, pp. 90–93, 2014.
- [24] A. Mackutė-Varoneckienė and T. Krilavičius, “Empirical study on unsupervised feature selection for document clustering.” *Human language technologies - the Baltic perspective : proceedings of the 6th international conference, Baltic HLT 2014*, pp. 107–110, 2014.
- [25] V. Pranckaitis and M. Lukoševičius, “Clustering of Lithuanian news articles,” in *CEUR Workshop Proceedings*, 2017.
- [26] J. Kapočiūtė-Dzikiėnė, R. Damaševičius, and M. Wozniak, “Sentiment analysis of Lithuanian texts using traditional and deep learning approaches.” *Computers*, vol. 8, no. 1:4, pp. 1–16, 2019.
- [27] J. Kapočiūtė-Dzikiėnė, R. Damaševičius, and M. Woźniak, “Sentiment analysis of Lithuanian texts using deep learning methods.” *Information and software technologies: 24th international conference, ICIST 2018, Vilnius, Lithuania, October 4–6, 2018: proceedings*, pp. 521–532, 2018.
- [28] J. Kapočiūtė-Dzikiėnė and R. Damaševičius, “Intrinsic evaluation of Lithuanian word embeddings using wordnet.” *Artificial intelligence and algorithms in intelligent systems: proceedings of 7th computer science on-line conference 2018*, pp. 394–404, 2018.
- [29] D. Chicco and G. Jurman, “The advantages of the matthews correlation coefficient (mcc) over f1 score and accuracy in binary classification evaluation,” *BMC genomics*, vol. 21, no. 1, p. 6, 2020.
- [30] A. Bielinskienė, L. Boizou, I. Bumbulienė, J. Kovalevskaitė, T. Krilavičius, J. Mandravickaitė, E. Rimkutė, and L. Vilkaitė-Lozdienė, “Lithuanian word embeddings,” 2019, CLARIN-LT digital library in the Republic of Lithuania. [Online]. Available: <http://hdl.handle.net/20.500.11821/26>