

Are Neural Language Models Good Plagiarists? A Benchmark for Neural Paraphrase Detection

Jan Philip Wahle
University of Wuppertal
Wuppertal, Germany
wahle@uni-wuppertal.de

Terry Ruas
University of Wuppertal
Wuppertal, Germany
ruas@uni-wuppertal.de

Norman Meuschke
University of Wuppertal
Wuppertal, Germany
meuschke@uni-wuppertal.de

Bela Gipp
University of Wuppertal
Wuppertal, Germany
gipp@uni-wuppertal.de

Abstract—The rise of language models such as BERT allows for high-quality text paraphrasing. This is a problem to academic integrity, as it is difficult to differentiate between original and machine-generated content. We propose a benchmark consisting of paraphrased articles using recent language models relying on the Transformer architecture. Our contribution fosters future research of paraphrase detection systems as it offers a large collection of aligned original and paraphrased documents, a study regarding its structure, classification experiments with state-of-the-art systems, and we make our findings publicly available.

Index Terms—Paraphrase detection, BERT, transformers

I. INTRODUCTION

Transformer-based language models [1] have reshaped Natural Language Processing (NLP) and became the standard paradigm for most of its downstream tasks [2], [3]. Now, Transformer-based models are ascending to other domains such as Computer Vision [4]. We anticipate these models will similarly influence Plagiarism Detection in the future. Plagiarism is characterized by the use of ideas, concepts, words, or structures without proper source acknowledgment, and often uses paraphrasing to conceal such practices [5]. Paraphrasing tools (e.g., *SpinBot*¹ and *SpinnerChief*²) help to perpetuate and increase the output of plagiarised content, challenging current Paraphrase Detection Systems (PDS).

Paraphrase tools and deterministic approaches used in machine-paraphrasing will soon give space to neural language models, which are already able to incorporate intrinsic features from human language effectively [3], [6]. The ability of models such as GPT-3 [3] to produce high-quality texts, similar to humans, raises an important concern in the Plagiarism Detection community, as statistical and traditional machine learning solutions are often not robust enough to distinguish semantically similar texts [7]. Classification based on Transformer models seem to be a natural choice to handle this new form of plagiarism. However, these solutions typically require sufficient labeled datasets during training for proper classification. As the use of neural language models for paraphrasing is a recent trend³, the lack of data to train PDS is amplified.

In this paper, we contribute towards future solutions in Plagiarism Detection of paraphrased text by publishing an extensive dataset collection⁴ of $\approx 3M$ paraphrased documents using neural language models, i.e. documents paraphrased by Transformers. To the best of our knowledge, this is the first dataset of its kind. Our dataset is derived from the original human written content of Wikipedia articles [7], scientific papers from arXiv, and Thesis documents [8]. To parse the 163 735 original paragraphs into their paraphrased version we use three state-of-the-art language models, i.e., BERT [2], RoBERTa [9], and Longformer [10]. We provide a comparison between original and paraphrased content regarding their paragraph embedding, showing our proposed dataset keeps the original meaning of its sources. We study how word-embeddings and the three Transformer models perform in classifying paraphrased documents to underline the difficulty of the proposed dataset. All data⁴ and source code⁵ is publicly available to collaborate with other researchers in the domain.

II. RELATED WORK

Paraphrase detection is widely explored in the domains of NLP and Digital Libraries [5]. To identify similar content between different elements, many approaches combine one or more techniques in lexical, syntactical, or semantic text analysis [7]. Among the resources used to validate PDS, PAN⁶ is arguably the most disseminated one, with challenges dating back to 2009 [11]. PAN comprehends a group of events (e.g., workshops) and tasks (e.g., authorship analysis) on digital text forensics and stylometry used as a benchmark for PDS. Another popular resource is the MRPC [12] dataset, a collection of human-annotated unique sentence pairs from news articles. The sentence pairs in MRPC capture paraphrasing and semantic equivalence relationships. Unfortunately, the Plagiarism Detection tasks offered in PAN and MRPC do not cover machine-generated cases, neither from paid methods (e.g., SpinBot) nor Transformer-based models (e.g., BERT).

GLUE [13] is one of the most popular benchmarks in NLP and contains the *Quora Question Pairs* task used to validate PDS. In this paraphrase identification task, systems need to

¹<https://spinbot.com>

²<http://www.spinnerchief.com/>

³<https://huggingface.co/models?search=paraphrase>

⁴<https://doi.org/10.5281/zenodo.4621403>

⁵Its access is granted upon the acceptance of its terms and conditions to avoid its misuse. Please visit our repository for more information.

⁶<https://pan.webis.de/>

classify duplicate questions with the same intent in Quora⁷, a knowledge repository where users share questions and answers about diverse topics. As in the PAN benchmark, GLUE does not account for machine-paraphrased content. Another aspect not covered in PAN and GLUE is the length of documents being evaluated. In all benchmarks, the text portions are relatively small (i.e., short sentences), which makes their generalization for Plagiarism Detection challenging.

In [7], [8], they try to mitigate the lack of machine-paraphrased datasets with a large collection of documents from arXiv, Thesis, and Wikipedia using paid paraphrasing tools (i.e., SpinBot and SpinnerChief). Their dataset is composed of paragraphs, promoting a more realistic plagiarism scenario [14]. Nevertheless, at the time of writing, no paraphrase detection benchmark makes use of Transformer-based architectures [1] to generate documents. Currently, the HuggingFace API³ offers few neural language models able to paraphrase text excerpts. Yet, most models are based on the same technique (i.e., T5 [6]), only work for short sentences, and do not offer a comprehensive dataset. In an attempt to mitigate the mentioned deficiencies, we release a new dataset⁴ composed of original paragraphs from arXiv, Thesis, and Wikipedia [7], [8], and their respective neural machine-paraphrased aligned pairs. We aim to foster future research in PDS that will soon need to compete against state-of-the-art NLP approaches used for text paraphrasing.

III. METHODOLOGY

Our machine-paraphrased corpus is derived from human-written English *featured articles*⁸ of Wikipedia from Foltynek et al. [7], arXiv scientific papers randomly sampled from the *no problems* category of the arXMLiv⁹ project, and randomly selected graduation Thesis of *English as a Second Language* (ESL) students at the Mendel University in Brno, Czech Republic from Wahle et al. [8]¹⁰.

Table I shows the details of the original content of all data sources considered. Similarly to [8], our training examples are derived only from Wikipedia (Wiki-Train) and the remaining data sources are used as test cases. Wikipedia provides a large collection for training, while other data sources are important to verify a model’s generalization. All documents are split at a paragraph level, as it reflects a more realistic scenario in the plagiarism domain compared to sentences [5], [14]. We use BERT [2], RoBERTa [9], and Longformer [10] to generate the neural paraphrased content for the entire dataset. BERT offers a strong baseline for classification, RoBERTa and Longformer improve BERT’s architecture through more training volume and an efficient attention mechanism, respectively.

To generate the dataset, we first exclude named entities, punctuation, brackets, digits, currency, and quote elements from paraphrasing to avoid producing false information, or inconsistent punctuation compared to their original source.

⁷<https://www.quora.com/about>

⁸https://en.wikipedia.org/wiki/Wikipedia:Content_assessment

⁹<https://kwarc.info/projects/arXMLiv/>

¹⁰<http://doi.org/10.5281/zenodo.3608000>

Next, we mask a certain percentage of words (cf. Section IV), forward it through each model, and obtain word candidates and their confidences. We choose the candidate with the highest confidence and replace the masked original words. We experimented with sampling uniformly over the top-k word predictions but neglected this method because of poor paraphrasing quality.

To identify whether our classification benchmark is challenging, we classify paraphrased and original paragraphs with four models. Wahle et al. [8] shows text-matching software fails to detect machine-paraphrased text while word embeddings, machine-learning classifiers, and particularly Transformer-based models, appear to generalize considerably well. Therefore, we chose the three BERT-related models used for paraphrasing, and a machine learning baseline composed of a Support Vector Machine (SVM) using features from fastText [15] representing traditional (sub-)word embeddings for classification. We limit the number of input tokens for each model to 512 for a fair comparison of the models without losing important context information¹¹. To generate the paraphrased training set (Wiki-Train) we used BERT, while for the paraphrased test set we generate one set with each of the three language models.

We publish the complete dataset including all model variations for generating the training and test sets in a public repository⁴ to enable researchers to extend our experiments. Evaluating each model with their own paraphrased text allows us to verify an assumption from related work, i.e., the best classifier is the language model used to generate paraphrased text [16]. We leave the investigation of additional neural language models and paraphrased variations to future work.

IV. ANALYSIS

We investigate the main components for our proposed neural paraphrased dataset. First, we perform an ablation study to understand how the Masked Language Model (MLM) probability affects the difficulty of our benchmark. Second, we evaluate how similar the generated paraphrased content regarding its original aligned pair. Lastly, we compare a SVM model using fastText [15] as features with three neural language models in a classification task. If not specified, all hyperparameters are used in their default configuration.

The Transformer-based models we selected, i.e., BERT, RoBERTa, and Longformer, are pre-trained in the MLM task. MLM masks a portion of words and the model has to infer the most probable word-choices for these unseen words. Considering each language model might predict different words, we evaluate how different masking probabilities affect the classification of paraphrased text in an ablation study. As Fig. 1 shows, our focus is to understand the role of MLM variations, so we generate all train-test pairs for each neural language model separately, and use the same feature-classifier to distinguish between original and paraphrased content. To encode all text as features, we used the sentence embedding

¹¹99.35% of the datasets’ text can be represented with less than 512 tokens.

TABLE I

OVERVIEW OF THE TOP CATEGORIES FOR NER AND POS TAGS FROM ORIGINAL CONTENT (TEST SETS) OF ARXIV, THESES AND WIKIPEDIA (TEST AND TRAIN SETS RESPECTIVELY) DATASETS. THE VALUES ARE ORDERED ACCORDING TO ARXIV. OTHERS CONTAINS ALL TAGS NOT DISPLAYED. FOR A FULL VERSION OF THE TABLE AND COMPLETE INFORMATION PLEASE VISIT OUR REPOSITORY⁴.

		arXiv	Thesis	Wiki	Wiki-Train
General	Paragraphs	20966	5226	39241	98282
	# Words	3194695	747545	5993461	17390048
	Avg. Words	152.38	143.04	152.73	176.94
POS	NOUN	22.22%	23.34%	17.36%	17.74%
	VERB	12.29%	12.81%	12.72%	12.96%
	ADP	12.28%	11.33%	12.11%	12.14%
	DET	12.03%	10.14%	10.08%	10.13%
	PUNCT	11.58%	11.80%	13.61%	13.31%
	ADJ	9.14%	7.86%	6.10%	6.53%
	ADV	3.83%	3.52%	2.96%	3.01%
	PROPN	3.74%	4.88%	12.35%	11.43%
	OTHERS	12.90%	14.32%	12.71%	12.76%
	Total	100.00%	100.00%	100.00%	100.00%
NER	CARDINAL	1.27%	1.12%	1.04%	0.92%
	ORG	0.98%	1.12%	1.73%	1.61%
	PERSON	0.48%	0.42%	2.13%	2.03%
	DATE	0.26%	1.00%	1.56%	1.51%
	GPE	0.22%	0.67%	1.17%	1.16%
	ORDINAL	0.18%	0.19%	0.37%	0.31%
	NORP	0.10%	0.27%	0.51%	0.59%
	PERCENT	0.05%	0.23%	0.05%	0.04%
	OTHERS	0.21%	0.27%	1.12%	1.03%
	Total	3.76%	5.29%	9.68%	9.20%

of fastText¹² (subword) trained on the full English Wikipedia from 2017, the UMBC webbase corpus and StatMT news dataset with 300 dimensions, and feed them into a SVM¹³ classifier. Fig. 1 shows higher MLM probabilities reduce the difficulty of the classification task. The correlation between the frequency of paraphrased content and F1-score is also verified in non-neural paraphrasing tools [8].

Commonly between all paraphrasing models, Thesis documents pose the most challenging scenario for fastText + SVM while arXiv and Wikipedia articles obtain higher classification performance which we assume is due to their similarity to the training set consisting only of Wikipedia articles. Factors like sub-optimal word choice and grammatical errors may increase the difficulty to classify Thesis documents correctly.

Masking only 15% of the content seems to provide the most challenging scenario for classification, and is the same masking probability used in pre-training BERT [2], thus, we use this configuration for the classification task in Table II. Paid online paraphrasing tools replace a similar proportion of words ranging from 12.58% to 19.37% words on average, supporting our choice of the masking probability parameter [8].

To evaluate the semantic representation of original and paraphrased test sets, we sample 30 paragraphs from all data sources. Fig. 2 (generated through t-SNE¹⁴) shows the neural paraphrased content overlaps significantly with their original

¹²<https://fasttext.cc/docs/en/english-vectors.html>

¹³<https://scikit-learn.org/stable/modules/svm.html>

¹⁴<https://scikit-learn.org/stable/modules/generated/sklearn.manifold.TSNE.html>

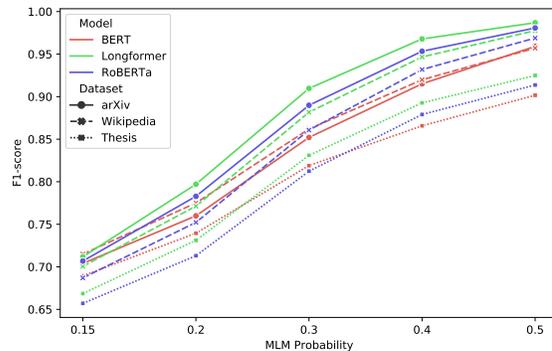


Fig. 1. F1-macro score of fastText + SVM for neural language models and masked language model probabilities [0.15-0.50].

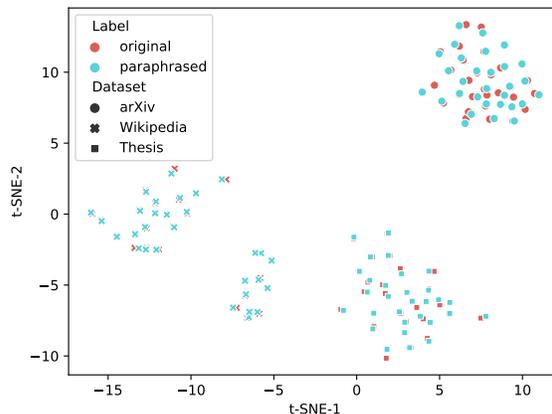


Fig. 2. BERT embeddings for paragraphs in a two-dimensional space reduced by the t-SNE algorithm. The embeddings of paraphrased documents are close to the original ones although 15% of the words have been changed.

counter-part although we attempt to change 15% of words. Thus, we can assume the Transformed-based models were able to keep the original texts’ semantics.

In the classification experiments (Table II), we want to explore the generalization aspects of each model regarding its ability to identify paraphrased content from other neural language models. We paraphrase Wiki-Train (Table I) using BERT and all test sets using each language model. The classification models are provided with mutually exclusive documents to avoid memorizing the differences between aligned documents. As each neural language model has its own paraphrased test set, we can investigate the findings of [16], which states the best classifier is the language model used to generate the paraphrased text.

Table II reports the F1-Micro score for each classification model on the paraphrased test sets for arXiv, Wikipedia, and Thesis documents. The baseline model (fastText + SVM) performed similarly among different paraphrasing models with F1=68.36% (RoBERTa) to F1=70.28% (BERT). Neural lan-

TABLE II
OVERVIEW OF F1-MACRO CLASSIFICATION RESULTS. THE HIGHEST RESULTS PER COLUMN ARE IN **BOLDFACE**.

Classification Models	Datasets	Paraphrased Models		
		BERT	RoBERTa	Longformer
fastText + SVM (baseline)	arXiv	70.40%	70.68%	71.17%
	Wikipedia	71.50%	68.70%	70.05%
	Thesis	68.94%	65.70%	66.85%
	Average	70.28%	68.36%	69.36%
BERT	arXiv	80.83%	68.90%	68.49%
	Wikipedia	83.21%	68.85%	69.46%
	Thesis	74.74%	67.39%	66.04%
	Average	79.59%	68.38%	68.00%
RoBERTa	arXiv	70.41%	85.40%	82.95%
	Wikipedia	72.18%	84.20%	82.15%
	Thesis	68.99%	79.13%	77.76%
	Average	70.53%	82.91%	80.95%
Longformer	arXiv	65.18%	85.46%	89.93%
	Wikipedia	69.98%	81.76%	86.03%
	Thesis	65.72%	77.96%	81.31%
	Average	66.96%	81.73%	85.76%

guage models consistently identified their own paraphrased text best, supporting the findings of [16] with F1=79.59% (BERT) to F1=85.76% (Longformer).

When we applied classification models to paraphrasing from other models (e.g., BERT classifies text that was paraphrased by Longformer), the classification models outperformed static word embeddings generally or achieved comparable scores. The overall scores for unseen paraphrasing models indicate the machine-paraphrased paragraphs construct a challenging benchmark. The scores for unseen paraphrasing models range from F1=68.00% to F1=81.73% with an average of 72.75% which is lower than for paid paraphrasing services, underlining the difficulty of this benchmark [8].

RoBERTa and Longformer presented similar results for all datasets, which we assume are because of their overlapping pre-training datasets. BERT uses a subset of RoBERTa and Longformer’s training data and identifies their paraphrased text with comparable F1-scores, supporting this hypothesis. RoBERTa achieved the best result on average for all paraphrasing techniques (F1=78.15%) positioning it as the most general model we tested for detecting neural machine-paraphrasing.

All classification models consistently identified Wikipedia articles the best, an expected outcome given the composition of our training corpus (i.e., Wikipedia). The three models identified arXiv articles similarly well due to their similarity to the training set, corroborating with our ablation study (cf. Fig. 1). Thesis documents from ESL students presented the most prominent challenge for our classification models which we presume results from their higher ratios of grammatical and linguistic errors.

V. FINAL CONSIDERATION AND FUTURE WORK

In this paper, we proposed a neural machine-paraphrased dataset⁴ with $\approx 3M$ paragraphs using three Transformer-based models, fostering future research in PDS. To the best of our

knowledge, this is the first time BERT [2], RoBERTa [9], and Longformer [10] are used to generate machine-paraphrased documents from Wikipedia, arXiv, and Thesis documents. We showed how similar neural paraphrased documents are in comparison to their original content, reinforcing the quality of neural paraphrase techniques. Furthermore, we showed Transformers are more effective in classifying original and paraphrased content when compared to static word embeddings (i.e., fastText) and Transformers are most effective when identifying their own paraphrased text. Our results show **RoBERTa** performed the best in detecting paraphrasing on average over all paraphrasing models. We leave to future work the investigation of auto-regressive models such as T5 [6] and GPT-3 [3].

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 30*. Curran Associates, Inc., 2017, pp. 5998–6008, <https://arxiv.org/abs/1706.03762>. [Online]. Available: <http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>
- [2] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” *arXiv:1810.04805 [cs]*, May 2019, arXiv: 1810.04805. [Online]. Available: <http://arxiv.org/abs/1810.04805>
- [3] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, and D. ... Amodei, “Language Models are Few-Shot Learners,” *arXiv:2005.14165 [cs]*, Jun. 2020, tex.ids: BrownMRS20a arXiv: 2005.14165. [Online]. Available: <http://arxiv.org/abs/2005.14165>
- [4] S. Khan, M. Naseer, M. Hayat, S. W. Zamir, F. S. Khan, and M. Shah, “Transformers in Vision: A Survey,” *arXiv:2101.01169 [cs]*, Feb. 2021, arXiv: 2101.01169. [Online]. Available: <http://arxiv.org/abs/2101.01169>
- [5] T. Foltýnek, N. Meuschke, and B. Gipp, “Academic Plagiarism Detection: A Systematic Literature Review,” *ACM Computing Surveys*, vol. 52, no. 6, pp. 112:1–112:42, 2019.
- [6] 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,” *ArXiv191010683 Cs Stat*, Jul. 2020.
- [7] T. Foltýnek, T. Ruas, P. Scharpf, N. Meuschke, M. Schubotz, W. Grosky, and B. Gipp, “Detecting Machine-obfuscated Plagiarism,” in *Proceedings of the iConference 2020*, ser. LNCS. Springer, 2020.
- [8] J. P. Wahle, T. Ruas, T. Foltýnek, N. Meuschke, and B. Gipp, “Identifying Machine-Paraphrased Plagiarism,” *arXiv:2103.11909 [cs]*, Jan. 2021, submitted. [Online]. Available: <http://arxiv.org/abs/2103.11909>
- [9] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “RoBERTa: A Robustly Optimized BERT Pretraining Approach,” *arXiv:1907.11692 [cs]*, Jul. 2019, arXiv: 1907.11692. [Online]. Available: <http://arxiv.org/abs/1907.11692>
- [10] I. Beltagy, M. E. Peters, and A. Cohan, “Longformer: The Long-Document Transformer,” *arXiv:2004.05150 [cs]*, Apr. 2020, arXiv: 2004.05150. [Online]. Available: <http://arxiv.org/abs/2004.05150>
- [11] M. Pothast, B. Stein, A. Eiselt, A. Barrón Cedeño, and P. Rosso, “Overview of the 1st International Competition on Plagiarism Detection,” in *Proceedings of the 3rd Workshop on Uncovering Plagiarism, Authorship and Social Software Misuse and 1st International Competition on Plagiarism Detection*, vol. 502 CEUR WS, 2009, pp. 1–9. [Online]. Available: <http://ceur-ws.org/Vol-502/paper1.pdf>
- [12] B. Dolan and C. Brockett, “Automatically constructing a corpus of sentential paraphrases,” in *Third International Workshop on Paraphrasing (IWP2005)*. Asia Fed. of Natural Language Processing, January 2005.
- [13] 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:1804.07461 [cs]*, Feb. 2019.
- [14] A. M. Rogerson and G. McCarthy, “Using Internet based paraphrasing tools: Original work, patchwriting or facilitated plagiarism?” *International Journal for Educational Integrity*, vol. 13, no. 1, 2017.

- [15] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, "Enriching Word Vectors with Subword Information," *Transactions of the Association for Computational Linguistics*, vol. 5, pp. 135–146, Dec. 2017.
- [16] R. Zellers, A. Holtzman, H. Rashkin, Y. Bisk, A. Farhadi, F. Roesner, and Y. Choi, "Defending Against Neural Fake News," *arXiv:1905.12616 [cs]*, Oct. 2019, arXiv: 1905.12616. [Online]. Available: <http://arxiv.org/abs/1905.12616>