

Identifying Machine-Paraphrased Plagiarism

Jan Philip Wahle^{1,2}, Terry Ruas³, Tomas Foltýnek³, Norman Meuschke³, and Bela Gipp³

¹University of Wuppertal

²Mendel University in Brno

³Affiliation not available

February 1, 2023

Abstract

Employing paraphrasing tools to conceal plagiarized text is a severe threat to academic integrity. To enable the detection of machine-paraphrased text, we evaluate the effectiveness of five pre-trained word embedding models combined with machine learning classifiers and state-of-the-art neural language models. We analyze preprints of research papers, graduation theses, and Wikipedia articles, which we paraphrased using different configurations of the tools SpinBot and SpinnerChief. The best performing technique, Longformer, achieved an average F1 score of 80.99% (F1=99.68% for SpinBot and F1=71.64% for Spinner-Chief cases), while human evaluators achieved F1=78.4% for SpinBot and F1=65.6% for SpinnerChief cases. We show that the automated classification alleviates shortcomings of widely-used text-matching systems, such as Turnitin and PlagScan. To facilitate future research, all data ³, code ⁴, and two web applications ⁵⁶ showcasing our contributions are openly available.

Related papers at <https://jpwahle.com/pub/> and <https://gipp.com/pub/>

J. P. Wahle, T. Ruas, T. Foltýnek, N. Meuschke, and B. Gipp (2022). Identifying Machine-Paraphrased Plagiarism. In: Smits, M. (eds) Information for a Better World: Shaping the Global Future. iConference 2022. Lecture Notes in Computer Science, vol 13192. Springer, Cham. doi: 10.1007/978-3-030-96957-8_34

Click to download: [BibTeX](#) | [RIS](#) | [ENW](#)

Identifying Machine-Paraphrased Plagiarism

Jan Philip Wahle¹[0000-0002-2116-9767], Terry Ruas¹[0000-0002-9440-780X],
Tomáš Foltýnek²[0000-0001-8412-5553], Norman Meuschke¹[0000-0003-4648-8198],
and Bela Gipp¹[0000-0001-6522-3019]

¹ University of Wuppertal, Rainer-Gruenter-Straße, 42119 Wuppertal, Germany
last@uni-wuppertal.de

² Mendel University in Brno, 61300 Brno, Czechia
tomas.foltynek@mendelu.cz

Abstract. Employing paraphrasing tools to conceal plagiarized text is a severe threat to academic integrity. To enable the detection of machine-paraphrased text, we evaluate the effectiveness of five pre-trained word embedding models combined with machine learning classifiers and state-of-the-art neural language models. We analyze preprints of research papers, graduation theses, and Wikipedia articles, which we paraphrased using different configurations of the tools SpinBot and SpinnerChief. The best performing technique, Longformer, achieved an average F1 score of 80.99% (F1=99.68% for SpinBot and F1=71.64% for SpinnerChief cases), while human evaluators achieved F1=78.4% for SpinBot and F1=65.6% for SpinnerChief cases. We show that the automated classification alleviates shortcomings of widely-used text-matching systems, such as Turnitin and PlagScan. To facilitate future research, all data³, code⁴, and two web applications^{5,6} showcasing our contributions are openly available.

Keywords: Paraphrase detection · plagiarism · document classification · transformers · BERT · Wikipedia

1 Introduction

Plagiarism is a pressing problem for educational and research institutions, publishers, and funding agencies [12]. To counteract plagiarism, many institutions employ *text-matching software*. These tools reliably identify duplicated text yet are significantly less effective for paraphrases, translations, and other concealed forms of plagiarism [11, 12].

Studies show that an alarming proportion of students employ *online paraphrasing tools* to disguise text taken from other sources [38, 40]. These tools employ artificial intelligence approaches to change text, e.g., by replacing words

³<https://doi.org/10.5281/zenodo.3608000>

⁴<https://github.com/jpelhaW/ParaphraseDetection>

⁵<http://purl.org/spindetector>

⁶<https://huggingface.co/jpelhaw/longformer-base-plagiarism-detection>

with their synonyms [56]. Paraphrasing tools serve to alter the content so that search engines do not recognize the fraudulent websites as duplicates.

In academia, paraphrasing tools help to mask plagiarism, facilitate collusion, and help ghostwriters with producing work that appears original. These tools severely threaten the effectiveness of text-matching software, which is a crucial support tool for ensuring academic integrity. The academic integrity community calls for technical solutions to identify the machine-paraphrased text as one measure to counteract paraphrasing tools [40].

The International Journal for Educational Integrity recently devoted a special issue⁷ to this topic.

We address this challenge by devising an automated approach that reliably distinguishes human-written from machine-paraphrased text and providing the solution as a free and open-source web application.

In this paper, we extend Foltýnek et al. [13] work by proposing two new collections created from research papers on arXiv⁸ and graduation theses of “English language learners” (ELL), and explore a second paraphrasing tool for generating obfuscated samples. We also include eight neural language models based on the Transformer architecture for identifying machine-paraphrases.

2 Related Work

The research on plagiarism detection technology has yielded many approaches that employ lexical, syntactical, semantic, or cross-lingual text analysis [12]. These approaches reliably find copied and moderately altered text; some can also identify paraphrased and machine-translated text. Methods to complement text analysis focus on non-textual features [27], such as academic citations [29], images [28], and mathematical content [30], to improve the detection of concealed plagiarism.

Most research on paraphrase identification quantifies to which degree the meaning of two sentences is identical. Approaches for this task employ lexical, syntactic, and semantic analysis (e.g., word embedding) as well as machine learning and deep learning techniques [12, 50].

The research on distinguishing machine-paraphrased text passages from original content is still in an early stage. Zhang et al. [56] provided a tool that determines if two articles are derived from each other. However, they did not investigate the task of distinguishing original and machine-fabricated text. Dey et al. [9] applied a Support Vector Machine (SVM) classifier to identify semantically similar tweets and other short texts. A very recent work studied word embedding models for paraphrase sentence pairs with word reordering and synonym substitution [1]. In this work, we focus on detecting paraphrases without access to pairs as it represents a realistic scenario without pair information.

Techniques to accomplish the task of paraphrase detection, dense vector representations of words in documents have attracted much research in recent

⁷<https://edintegrity.biomedcentral.com/mbp>

⁸<https://arxiv.org>

years. Word embedding techniques, such as word2vec [31], have alleviated common problems in bag-of-words (BOW) approaches, e.g., scalability issues and the curse of dimensionality. Representing entire documents in a single fixed-length dense vector (doc2vec) is another successful approach [23]. Word2vec and doc2vec can both capture latent semantic meaning from textual data using efficient neural network language models. Prediction-based word embedding models, such as word2vec and doc2vec, have proven themselves superior to count-based models, such as BOW, for several problems in Natural Language Processing (NLP), such as quantifying word similarity [42], classifying documents [41], and analyzing sentiment [36]. Gharavi et al. employed word embeddings to perform text alignment for sentences [14]. Hunt et al. integrated features from word2vec into machine learning models (e.g., logistic regression, SVM) to identify duplicate questions in the Quora dataset [16]. We, on the other hand, consider text documents generated with the help of automated tools at the paragraph level.

Recently, the NLP community adapted and extended the neural language model BERT [8] for a variety of tasks [2, 5, 33, 34, 45, 49, 54], similar to the way that word2vec [31] has influenced many later models in NLP [4, 41, 42]. Based on the Transformer architecture [48], BERT employs two pre-training tasks, i.e., *Masked Language Model* (MLM) and *Next Sentence Prediction* (NSP), to capture general aspects of language. MLM uses a deep bidirectional architecture to build a language model by masking random tokens in the input. The NSP task identifies if two sentences are semantically connected. The ALBERT [20], DistilBERT [44], and RoBERTa [25] models are all based on BERT and either improve their predecessor’s performance through hyperparameter adjustments or make BERT less computationally expensive. Different from ELMo [37] and GPT [39], BERT considers left-to-right and right-to-left context simultaneously, allowing a more realistic representation of the language. Although ELMo does use two LSTM networks, their weights are not shared during training. On top of MLM and NSP, BERT requires fine-tuning to specific tasks to adjust its weights accordingly.

Other recent models proposed architectural and training modifications for BERT. ELECTRA changes BERT’s MLM task to a generator-discriminator setup [5]. Tokens are substituted with artificially generated ones from a small masked language model and discriminated in a noise contrastive learning process [15]. BART pre-trains a bidirectional auto-encoding and an auto-regressive Transformer in a joint structure [24]. The two-stage denoising auto-encoder first corrupts the input with an arbitrary function (bidirectional) and uses a sequence-to-sequence approach to reconstruct the original input (auto-regressive) [24]. In XLNet, a permutation language modeling predicts one word given its preceding context at random [54]. Longformer proposed the most innovative contribution by exploring a new scheme for calculating attention [3]. Longformer’s attention mechanism combines windowed local with global self-attention while also scaling linearly with the sequence length compared to earlier models (e.g., RoBERTa).

Foltýnek et al. [13] tested the effectiveness of six word embedding models and five traditional machine learning classifiers for identifying machine-paraphrased.

We paraphrased Wikipedia articles using the SpinBot⁹ API, which is the technical backbone of several widely-used services, such as Paraphrasing Tool¹⁰ and Free Article Spinner¹¹ [40]. The limitations of [13] are the exclusive use of one data source, the lack of recent neural language models, and the reliance on a single paraphrasing tool. In this paper, we address all three shortcomings by considering arXiv and graduation theses as new data sources (Section 3.2), eight neural language models (Section 3.5), and SpinnerChief¹² as an additional paraphrasing tool (Section 3.1).

Lan et al. [19] compared five neural models (e.g., LSTM and CNN) using eight NLP datasets, of which three focus on sentence paraphrase detection (i.e., Quora [17], Twitter-URL [18], and PIT-2015 [53]). Subramanian et al. presented a model that combines language modeling, machine translation, constituency parsing, and natural language inference in a multi-task learning framework for sentence representation [46]. Their model produces state-of-the-art results for the MRPC [10] dataset. Our experiments consider a multi-source paragraph-level dataset and more recent neural models to reflect a real-world detection scenario and investigate recent NLP techniques that have not been investigated for this use case before.

Wahle et al. [50] is the only work, to date, that applies neural language models to generate machine paraphrased text. They use BERT and other popular neural language models to paraphrase an extensive collection of original content. We plan to investigate additional models and combine them with the work on generating paraphrased data [50], which could be used for training.

3 Methodology

Our primary research objective is to provide a free service that distinguishes human-written from machine-paraphrased text while being insensitive to the topic and type of documents and the paraphrasing tool used. We analyze paragraphs instead of sentences or entire documents since it represents a more realistic detection task [40, 52]. Sentences provide little context and can lead to more false positives when sentence structures are similar. Fulltext documents are computationally expensive to process, and in many cases the extended context does not provide a significant advantage over paragraphs. We extend Foltýnek et al.’s [13] study by analyzing two new datasets (arXiv and theses), including an extra machine-paraphrasing tool (SpinnerChief), and evaluating eight state-of-the-art neural language models based on Transformers [48]. We first performed preliminary experiments with classic machine learning approaches to identify the best-performing baseline methods for paraphrasing tools and datasets we investigate. Next, we compared the best-performing machine learning techniques to

⁹<https://spinbot.com/>

¹⁰<https://paraphrasing-tool.com/>

¹¹<https://free-article-spinner.com/>

¹²<http://www.spinnerchief.com/>

neural language models based on the Transformer architecture, representing the latest advancements in NLP.

3.1 Paraphrasing Tools

We employed two commercial paraphrasing services, i.e., *SpinBot*⁹ and *SpinnerChief*¹², to obfuscate samples in our training and test sets. We used SpinBot to generate the training and test sets and SpinnerChief only for the test sets.

SpinnerChief allows specifying the ratio of words it tries to change. We experimented with two configurations: the *default frequency* (*SpinnerChief-DF*), which attempts to change every fourth word, and an *increased frequency* (*SpinnerChief-IF*), which attempts to change every second word.

3.2 Datasets for Training and Testing

Most paraphrasing tools are paid services, which prevents experimenting with many of them. The financial costs and effort required for obtaining and incorporating tool-specific training data would be immense. Therefore, we employed transfer learning, i.e., used pre-trained word embedding models, trained the classifiers in our study on samples paraphrased using SpinBot, and tested whether the classification approach can also identify SpinnerChief’s paraphrased text.

Training Set: We reused the paragraph training set of Foltýnek et al. [13] and paraphrased all 4,012 *featured articles* from English Wikipedia using SpinBot⁹. We chose featured Wikipedia articles because they objectively cover a wide range of topics in great breadth and depth¹³. Approx. 0.1% of all Wikipedia articles carry the label *featured article*.

Thus, they are written in high-quality English by many authors and unlikely to be biased towards individual writing styles.

The training set comprises of 200,767 paragraphs (98,282 original, 102,485 paraphrased) extracted from 8,024 Wikipedia articles. We split each Wikipedia article into paragraphs and discarded those with fewer than three sentences, as Foltýnek et al. [13] showed such paragraphs often represent titles or irrelevant information.

Test Sets: Our study uses three test sets that we created from preprints of research papers on arXiv, graduation theses, and Wikipedia articles. Table 1 summarizes the test sets. For generating the *arXiv* test set, we randomly selected 944 documents from the *no problems* category of the arXMLiv project¹⁴. The *Wikipedia* test set is identical to [13]. The paragraphs in the test set were generated analogously to the training set. The *theses* test set comprises paragraphs in 50 randomly selected graduation theses of ELL at the Mendel University in

¹³https://en.wikipedia.org/wiki/Wikipedia:Content_assessment

¹⁴<https://kwarc.info/projects/arXMLiv/>

Brno, Czech Republic. The theses are from a wide range of disciplines, e.g., economics, computer science, and cover all academic levels. Unlike the arXiv and Wikipedia documents, the theses were only available as PDF files, thus required conversion to plain text. We removed all content before the introduction section of each thesis, the bibliography, and all appendices to avoid noisy data (e.g., table of contents).

Table 1. Overview of the test sets.

No. of paragraphs	arXiv		theses		Wikipedia	
	Original	Para-phrased	Original	Para-phrased	Original	Para-phrased
SpinBot	20,966	20,867	5,226	3,463	39,261	40,729
SpinnerChief-DF	20,966	21,719	2,379	2,941	39,261	39,697
SpinnerChief-IF	20,966	21,671	2,379	2,941	39,261	39,618

3.3 Word Embedding Models

Table 2 summarizes the word embedding models analyzed in our experiments: GloVe¹⁵ [35], word2vec¹⁶(w2v) [31], fastText¹⁷(FT-rw and FT-sw) [4], and doc2vec (d2v) [23] that we trained from scratch. The d2v model uses a distributed bag-of-words training objective, a window size of 15 words, a minimum count of five words, trained word-vectors in skip-gram fashion, averaged word vectors, and 30 epochs. All word embedding models have 300 dimensions. Parameters we do not explicitly mention correspond to the default values in the *gensim* API¹⁸.

Our rationale for choosing the pre-trained word embedding models was to explore the most prominent techniques regarding their suitability for the plagiarism detection task. GloVe [35] builds a co-occurrence matrix of the words in a corpus and explores the word probabilities ratio in a text to derive its semantic vectors as a count-based model. The training of w2v tries to predict a word given its context (cbow) or the context given a word (skip-gram) [31].

Even though numerous NLP tasks routinely apply GloVe and w2v [6, 41, 42], they do not consider two important linguistic characteristics: word ordering and sub-wording. To explore these characteristics, we also included fastText [4] and the paragraph vector model [23]. FastText builds its word representation by extending the skip-gram model with the sum of the n-grams of its constituent sub-word vectors. As the paraphrasing algorithms used by plagiarists are unknown, we hypothesize rare words can be better recognized by fastText through sub-wording. Two training options exist for the d2v model—Distributed Memory

¹⁵<https://nlp.stanford.edu/projects/glove/>

¹⁶<https://code.google.com/archive/p/word2vec/>

¹⁷<https://fasttext.cc/docs/en/english-vectors.html>

¹⁸<https://radimrehurek.com/gensim/models/doc2vec.html>

Model of Paragraph Vectors (pv-dm) and Distributed Bag of Words version of Paragraph Vector (pv-dbow). The former is akin to w2v cbow, while the latter is related to w2v skip-gram. Both options introduce a new paragraph-id vector that updates each context window on every timestamp. The paragraph-id vector seeks to capture the semantics of the embedded object. We chose a pv-dbow over a pv-dm model because of its superior results in semantic similarity tasks [22].

Table 2. Word embedding models in our experiments.

Algorithm	Main Characteristics	Training Corpus
GloVe	Word-word co-occurrence matrix	Wikipedia Dump 2014 + Gigaword 5
word2vec	Continuous Bag-of-Words	Google News
pv-dbow	Distributed Bag-of-Words	Wikipedia Dump 2010
fastText-rw	Skip-gram without sub-words	Wikipedia Dump 2017 + UMBC
fastText-sw	Skip-gram with sub-words	Wikipedia Dump 2017 + UMBC

In our experiments, we represented each text as the average of its constituent word vectors by applying the word embedding models in Table 2 [41, 43]. All models, except for d2v, yield a vector representation for each word. D2v produces one vector representation per document. Inferring the vector representations for unseen texts requires an additional training step with specific parameter tuning. We performed this extra training step with hyperparameters according to [41] for the *gensim* API: $\alpha = 10^{-4}$, $\alpha_{min} = 10^{-6}$, and 300 epochs [22]. The resulting pv-dbow embedding model requires at least 7 GB of RAM, compared to 1-3 GB required for other models. The higher memory consumption of pv-dbow can make it unsuitable for some use cases.

3.4 Machine Learning Classifiers

After applying the pre-trained models to our training and test sets, we passed on the results to three machine learning classifiers: Logistic Regression (LR), Support Vector Machine (SVM), and Naïve Bayes (NB).

We employed a grid-search approach implemented using the scikit-learn package¹⁹ in Python for finding the optimal parameter values for each classifier (Table 3)

3.5 Neural Language Models

We investigate the following neural language models based on the Transformer architecture: BERT [7], RoBERTa [26], ALBERT [21], DistilBERT [44], ELECTRA [5], BART [24], XLNet [54], and Longformer [3]. Our rationale for testing Transformer-based models is their ability to generally outperform traditional

¹⁹<https://scikit-learn.org>

Table 3. Grid-search parameters for ML classifiers.

Classifier	Parameter	Range
Logistic Regression	solver	newton-cg, lbfgs, sag, saga
	maximum iteration	500, 1000, 1500
	multi-class tolerance	ovr, multinomial 0.01, 0.001, 0.0001, 0.00001
Support Vector Machine	kernel	linear, radial bases function, polynomial
	gamma	0.01, 0.001, 0.0001, 0.0001
	polynomial degree C	1, 2, 3, 4, 5, 6, 7, 8, 9 1, 10, 100

word embedding and machine learning models in similar NLP tasks (e.g., document similarity). We chose the aforementioned models specifically because of two reasons. First, we explore models closely related or based on BERT that improve BERT through additional training time and data (RoBERTa) or compress BERT’s architecture with minimal performance loss (DistilBERT, ALBERT). Second, we use contrasting models to BERT that, although relying on the Transformer architecture, significantly change the training objective (XLNet), the underlying attention mechanism (Longformer), or employ a discriminative learning approach (ELECTRA, BART).

To classify whether a paragraph is paraphrased, we attach a randomly initialized linear layer on top of the model’s embedding of the aggregate token (e.g., [CLS] for BERT) to transform the embedding into binary space. The final layer predicts whether a paragraph has been paraphrased using cross-entropy loss. For all models, we use the base version, the official pre-trained weights, and the following configurations: a sequence length of 512 tokens, an accumulated batch size of 32, the Adam optimizer with $\alpha = 2e - 5$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e - 8$, and PyTorch’s native automated mixed-precision format. Using a common sequence length of 512 tokens allows for a fair comparison of the models without losing important context information²⁰. Section 4.1 provides more details about the models’ characteristics.

4 Evaluation

To quantify the effectiveness of classification approaches in identifying machine-paraphrased text, we performed three experiments. Section 4.1 presents the results of applying the pre-trained word embedding models in combination with machine learning classifiers and neural language models to the three test sets. Sections 4.2 and 4.3 indicate how well human experts and a text-matching software identify machine-paraphrased text to put the results of automated classification approaches into context.

²⁰99.35% of the datasets’ text can be represented with less than 512 tokens.

4.1 Automated Classification

This section presents the micro-averaged F1 scores (F1-Micro) for identifying paragraphs we paraphrased using either SpinBot or SpinnerChief and classified using combinations of pre-trained word embedding models and machine learning classifiers or Transformer-based language models.

Results of ML Techniques for SpinBot: Section 4.1 shows the results for classifying SpinBot test sets derived from arXiv, theses, and Wikipedia using combinations of pre-trained word embedding models and machine learning classifiers. GloVe, in combination with SVM, achieved the best classification performance for all test sets. The combination of w2v and SVM performed nearly as well as GloVe+SVM for all test sets. For the theses and Wikipedia test sets, the performance difference between GloVe+SVM and w2v+SVM is less than 2%, and for the arXiv test set 6.66%. All pre-trained word embedding models achieved their best results for the Wikipedia test set.

Table 4. Classification results (F1-Micro) of ML techniques for SpinBot.

	GloVe	w2v	d2v	FT-rw	FT-sw
arXiv	86.46	79.80	72.40	78.40	74.14
LR	76.53	74.82	69.42	75.08	65.92
SVM	86.46	79.80	72.40	76.31	74.15
NB	79.17	74.23	57.99	78.40	64.96
theses	83.51	81.94	61.92	72.75	64.78
LR	68.55	72.89	59.97	69.17	64.03
SVM	83.51	81.94	61.92	72.75	64.78
NB	75.22	74.18	42.30	72.11	61.99
Wikipedia	89.55	87.27	83.04	86.15	82.57
LR	80.89	84.50	81.08	85.13	78.97
SVM	89.55	87.27	83.04	86.15	82.57
NB	69.68	69.84	58.88	70.05	64.47

➤ **Boldface** indicates the best score for each test set, i.e., arXiv, theses, and Wikipedia. The score of the best-performing combination of embedding model and classifier is repeated in the row of the test set.

All classification approaches, except for w2v+SVM, performed worst for the theses test set. However, the drop in performance for theses test cases is smaller than we expected. The F1-Micro score of the best approach for the theses test set (GloVe+SVM) is 6.04% lower than for the Wikipedia test set and 3.09% lower than for the arXiv test set. This finding suggests that the quality of writing in student theses mildly affects the detection of machine-paraphrased text.

Although d2v seeks to mitigate shortcomings of its predecessor w2v, such as ignoring word order and producing a variable-length encoding, w2v surpassed d2v for all test sets. A possible reason is the short length of the paragraphs we consider. Lau et al. found that d2v’s performance decreases for short documents [22]. The results for paragraphs in Section 4.1 and for documents in our preliminary study [13], where d2v was the best-performing approach, support this conclusion.

For fastText (FT-rw and FT-sw in Section 4.1), we observe the same behavior as for w2v and d2v. The sub-word embeddings of FT-sw should provide a benefit over FT-rw, which encodes whole words, by capturing sub-word structures [4]. Therefore, we expected a better performance of FT-sw compared to FT-rw. However, FT-rw and simpler models, i.e., GloVe and w2v, performed better than FT-sw for all test sets.

Results of ML Techniques for SpinnerChief: Table 5 shows the results of applying the pre-trained word embedding models and machine learning classifiers to the arXiv, theses, and Wikipedia test sets containing cases paraphrased by the SpinnerChief tool. We either used the tool’s default setting, i.e., attempting to replace every fourth word (SpinnerChief-DF), or increased the frequency of attempted word replacements to every other word (SpinnerChief-IF).

Table 5. Classification results (F1-Micro) of ML techniques for SpinnerChief.

	SpinnerChief-DF					SpinnerChief-IF				
	GloVe	w2v	d2v	FT-rw	FT-sw	GloVe	w2v	d2v	FT-rw	FT-sw
arXiv	58.48	59.78	56.46	57.42	59.72	64.34	65.89	59.27	63.70	63.66
LR	52.14	55.43	56.46	57.42	58.64	54.92	59.61	59.07	61.74	61.57
SVM	58.42	57.65	56.43	56.43	59.72	64.12	62.77	59.27	62.97	63.66
NB	58.48	59.78	51.58	51.58	55.21	64.34	65.89	52.21	63.70	59.33
theses	52.63	53.60	59.09	53.08	57.25	58.57	58.24	63.15	59.13	61.27
LR	48.42	53.60	59.09	52.51	55.63	52.08	57.94	62.88	59.13	60.65
SVM	52.63	51.54	59.00	53.08	57.25	58.57	57.78	63.15	58.12	61.27
NB	50.90	53.32	54.94	52.78	46.99	55.62	58.24	55.09	57.19	50.13
Wikipedia	57.86	60.30	55.99	59.19	59.62	64.16	66.83	60.94	65.35	66.41
LR	52.97	55.90	55.64	56.40	59.62	55.68	61.32	60.16	62.51	66.41
SVM	57.09	57.48	55.99	57.15	58.72	64.16	64.56	60.94	63.61	64.81
NB	57.86	60.30	51.64	59.19	57.29	63.46	66.83	52.64	65.35	62.06

➤ *-DF* default frequency, *-IF* increased frequency (attempt changing every fourth or every second word).

➤ **Boldface** indicates the best score for each test set. The score of the best-performing combination of embedding model and classifier is repeated in the row of the test set.

We observe a drop in the SpinnerChief’s classification performance compared to the results for SpinBot. The average decrease in the F1-Micro scores was approx. 17% when using SpinnerChief-DF and approx. 13% for -IF. The comparison between the results of SpinBot and SpinnerChief-IF is more informative than comparing SpinBot to SpinnerChief-DF since the IF setting yields a similar ratio of replaced words to SpinBot.

As in SpinBot, all approaches performed best for the Wikipedia and worst for the theses. However, the performance differences were smaller for SpinnerChief than for SpinBot. For all SpinnerChief (DF and IF), the lowest F1-Micro scores were at most 6.5% below the highest scores, and the runner-ups were generally within an interval of 2% of the best scores.

The characteristics of ELL texts, e.g., sub-optimal word choice and grammatical errors, decreased the classification performance less than we had expected. The highest scores for the SpinBot theses are approx. 6% lower than the highest scores for any other dataset for SpinBot. For SpinnerChief, this difference is approx. 2%.

Notably, SpinnerChief’s settings for a stronger text obfuscation (SpinnerChief-IF) increased the rate with which the classification approaches identified the paraphrases. On average, SpinnerChief-DF replaced 12.58% and SpinnerChief-IF 19.37% of the words in the text (Section 3.1). The 6.79% increase in the number of replaced words for SpinnerChief-IF compared to SpinnerChief-DF increased the average F1-Micro score of the classification approaches by 5.56%. This correlation suggests that the classification approaches can recognize most of the characteristic word replacements that paraphrasing tools perform.

Text-matching software, such as Turnitin and PlagScan, are currently the de-facto standard technical support tools for identifying plagiarism. However, since these tools search for identical text matches, their detection effectiveness decreases when the number of replaced words increases (Table 7). Including additional checks, such as the proposed classification approaches, as part of the text-matching software detection process could alleviate the weaknesses of current systems.

We attribute the drop in the classification performance and the overall leveling of the F1-Micro scores for SpinnerChief test sets compared to SpinBot test sets to our transfer learning approach. As explained in Section 3, we seek to provide a system that generalizes well for different document collections and paraphrasing tools. Therefore, we used the machine-paraphrased text samples of SpinBot and applied the pre-trained word embedding models from Table 2 to extract the vector representations. We then used these vectors as the features for the machine learning classifiers for both Spinbot and SpinnerChief test sets.

We selected the combinations of word embedding models and machine learning classifiers that performed best for SpinBot (Section 4.1) and SpinnerChief (Table 5) as the baseline to which we compare the Transformer-based language models in the following section.

Results for Transformer-based Language Models: Table 6 shows the classification results of neural language models applied to all SpinBot and SpinnerChief

Table 6. Classification results (F1-Micro) of best ML techniques and neural language models for SpinBot and SpinnerChief.

Techniques	SpinBot			SpinnerChief-DF			SpinnerChief-IF		
	arXiv	Theses	Wiki	arXiv	Theses	Wiki	arXiv	Theses	Wiki
Baseline	86.46 ^a	83.51 ^a	89.55 ^a	59.78 ^b	59.09 ^c	60.30 ^b	65.89 ^b	63.15 ^d	66.83 ^b
BERT	99.44	94.72	99.85	50.74	50.42	43.00	64.59	63.59	57.45
ALBERT	98.91	96.77	99.54	66.88	47.92	50.43	75.57	56.75	59.61
DistilBERT	99.32	96.61	99.42	38.37	45.07	37.05	47.25	51.44	46.81
RoBERTa	99.05	97.34	99.85	57.10	47.40	48.03	66.00	58.24	58.94
ELECTRA	99.20	96.85	99.41	43.83	44.95	56.30	60.77	63.11	75.92
BART	99.58	99.66	99.86	69.38	53.39	48.62	76.07	63.57	58.34
XLNet	99.65	98.33	99.48	69.90	53.06	50.51	80.56	71.75	61.83
Longformer	99.38	99.81	99.87	76.44	70.15	63.03	78.34	74.82	67.11

➤ ^aGloVe+SVM ^bw2v+NB ^cd2v+LR ^dd2v+SVM

➤ The first horizontal block shows the best results of machine learning techniques, the second of models that optimize BERT, and the third of models that use new architectural or training approaches.

➤ **Boldface** indicates the best score for each test set.

test sets. The machine learning technique that performed best for each test set (Section 4.1 and table 5) is shown as *Baseline*.

For the SpinBot, all Transformer-based models outperformed their machine learning counterparts on average by 16.10% for theses, 13.27% for arXiv, and 10.11% for Wikipedia. Several models consistently achieved F1-Micro scores above 99% for all SpinBot cases. These findings show that the models could capture the intrinsic characteristics of SpinBot’s paraphrasing method very well. We stopped the training for each model after one epoch to avoid overfitting.

All techniques performed worse for SpinnerChief than for SpinBot, which we expected given the transfer learning approach. The drop in the classification performance was consistently lower for SpinnerChief-IF, which exhibits a similar ratio of replaced words as SpinBot, than for SpinnerChief-DF, which contain fewer replaced words than SpinBot. The most significant improvements in the scores for SpinnerChief-IF over SpinnerChief-DF are 16.94% for arXiv (ELECTRA), 18.69% for the theses (XLNet), and 19.62% for Wikipedia (BART).

These results show that the ratio of replaced words is a significant indicator of a models’ performance. However, since the paraphrasing methods of SpinBot and SpinnerChief (DF and IF) are unknown and could be different for each setting, one can interpret this finding in two ways. First, the models may capture the frequency of replaced words intrinsically and increase their attention to more words, which would mean the models can better detect more strongly altered paragraphs, such as those produced by SpinnerChief-IF. Second, SpinnerChief-IF cases might be better detectable because the paraphrasing method associated

with the SpinnerChief-IF setting might be more akin to the one of SpinBot than the method associated with the SpinnerChief-DF setting.

For all SpinnerChief-DF cases, Longformer consistently achieved the best results, surpassing the F1-Micro scores of the machine learning baselines by 10.15% on average and 16.66% for arXiv. For SpinnerChief-IF, XLNet, Longformer, and ELECTRA achieved the best results with an improvement in the F1-Micro scores of 14.67%, 11.67%, and 9.09% over the baseline scores for the arXiv, theses, and Wikipedia, respectively. As ELECTRA was pre-trained using a Wikipedia dump and the Books Corpus [57], we assume it also captured semantic aspects of Wikipedia articles.

The larger diversity in the training data of Longformer and XLNet (i.e., Gigaword 5 [32], CC Stories [47], and Realnews [55]) seems to enable the models to capture unseen semantic structures in the arXiv and theses better than other models.

BERT and its derived models performed comparably to the baselines for most SpinnerChief cases. DistilBERT, which uses knowledge distillation to reduce the number of parameters by 40% compared to BERT, performed significantly worse than its base model. For the SpinnerChief, the F1-micro scores for DistilBERT were on average 10.63% lower than for BERT, often falling into a score range achievable by random guessing. Although we expected a slight decline in the accuracy of DistilBERT compared to BERT due to the parameter reduction, the results fell well below our predictions. In comparison, for the General Language Understanding (GLUE) dataset [51], DistilBERT performed only 2.5% worse than BERT. ALBERT’s parameter reduction techniques, e.g., factorized embedding parametrization and parameter sharing, seem to be more robust. ALBERT outperformed BERT on average by 4.56% on the SpinnerChief test sets. With an average improvement of 0.99%, RoBERTa performed slightly better than BERT. However, as RoBERTa uses more parameters than most other BERT-related models and has exceptionally high computational requirements for pre-training, we rate this performance benefit as negligible.

In summary, improvements of BERT’s attention mechanism or training objective outperformed other BERT-based models for the machine-paraphrase detection task. We hypothesize the windowed local and global self-attention scheme used in Longformer allowed the model to generalize better between different paraphrasing tools. In eight out of nine scenarios, Longformer was either the best or second-best model overall. Also, for almost all cases, the neural language models surpassed the machine learning approaches’ results, thus providing a better solution to the problem of identifying machine-paraphrases. We see the SpinnerChief-DF test results set as a lower bound regarding the detection effectiveness for unseen spinning methods, even if the frequency of word replacements is significantly different from the frequency in our training set.

4.2 Human Baseline

To complement the earlier study of Foltýnek [13], we conducted a user survey with excerpts from ten randomly selected Wikipedia articles. We paraphrased

three of the ten excerpts using SpinnerChief-DF, three others using SpinnerChief-IF, and kept four excerpts unaltered. Using QuizMaker²¹, we prepared a web-based quiz that showed the ten excerpts one at a time and asked the participants to vote whether the text had been machine-paraphrased. We shared the quiz via e-mail and a Facebook group with researchers from the academic integrity community and 32 participants joined our study.

The participants' accuracy ranged between 20% and 100%, with an average of 65.59%. Thus, the F1-Micro score of the 'average' human examiner matched the average of the best scores of automated classification approaches for the SpinnerChief-IF test sets (65.29%). Some participants pointed out that oddness in the text of some excerpts, e.g., lowercase letters in acronyms, helped them identify the excerpts as paraphrased. For SpinBot, 73 participants answered the survey with an accuracy ranging from 40% to 100% (avg. 78.40%) according to [13].

Our experiments show that experienced educators who read carefully and expect to encounter machine-paraphrased text could achieve an accuracy between 80% and 100%. However, even in this setting, the average accuracy was below 80% for SpinBot and below 70% for SpinnerChief. We expect that the efficiency will be lower in a realistic scenario, in which readers do not pay special attention to spotting machine paraphrases.

4.3 Text-matching Software Baseline

To quantify the benefit of our automated classification over text-matching software, we tested how accurately current text-matching tools identify paraphrased text. We tested two systems—Turnitin, which has the largest market share, and PlagScan, which was one of the best-performing systems in a comprehensive test conducted by the European Network for Academic Integrity (ENAI) [11]. Our main objective was to test the tools' effectiveness in identifying patchwriting, i.e., inappropriately paraphrasing copied text by performing minor changes and substitutions. Patchwriting is a frequent form of plagiarism, particularly among students.

For this test, we created four sets of 40 documents each (160 documents total). We composed each document by randomly choosing 20 paragraphs from Wikipedia articles (2x40 documents), arXiv preprints (40 documents), and theses (40 documents). For each set of 40 documents, we followed the same scheme regarding the length and obfuscation of the chosen paragraphs. First, we created ten documents by varying the paragraphs' length taken from the source from one to ten sentences. In addition to using the ten documents unaltered, we also paraphrased all ten documents using SpinBot, SpinnerChief-DF, and SpinnerChief-IF.

To ensure this test is objective and comparable across the data sets, we exclusively report the overall percentages of matching text reported by a system (Table 7). In most cases, the systems identified the correct source. However, the

²¹<https://www.quiz-maker.com/>

Table 7. Percentage of text overlap reported by the text-matching systems Turnitin and PlagScan

Detection Corpus		arXiv	Theses	Wikipedia
Turnitin	Original	84.0	5.4	98.7
	SpinBot	7.0	1.1	30.2
	SpinnerChief-DF	58.5	4.0	74.5
	SpinnerChief-IF	38.8	1.2	50.1
PlagScan	Original	44.6	22.3	65.0
	SpinBot	0.0	0.1	0.5
	SpinnerChief-DF	9.2	12.0	19.1
	SpinnerChief-IF	1.8	0.5	3.1

systems often reported false positives caused by random matches, which means the systems’ actual retrieval effectiveness is slightly lower than reported.

The results show that PlagScan struggled to identify patchwriting. Even though the system indexes Wikipedia and could identify entirely plagiarized documents in the ENAI test [11], the average percentage of matching text reported for our patch-written documents was 63%. Paraphrasing documents using SpinBot and SpinnerChief-IF consistently prevented PlagScan from identifying the plagiarism. The average reported percentage of matching text was only 1% for SpinBot and 3% for SpinnerChief-IF test cases. For SpinnerChief-DF test cases, PlagScan could identify 19% of plagiarism present in documents, likely due to the lower ratio of altered words. Nevertheless, we can conclude obfuscating patch-written documents using a machine-paraphrasing tool likely prevents PlagScan from identifying plagiarism.

As shown in Table 7, Turnitin performed better for patch-written documents than PlagScan. For Wikipedia test cases, Turnitin reported 100% matching text for almost all cases. The average percentage of matching text Turnitin reported for machine-paraphrased documents was much higher than for PlagScan—31% for SpinBot, 74% for SpinnerChief-DF, and 50% for SpinnerChief-IF. However, machine-paraphrasing still prevents Turnitin from identifying a significant portion of the plagiarized content. Notably, Turnitin appears to index fewer theses than PlagScan, thus failing to report suspiciously high percentages of matching text for any theses test set, including ones containing unaltered paragraphs copied from theses.

For both systems, we observed the longer a plagiarized passage is, the more likely text-matching tools identified it. This result corresponds to the results of our classification approaches, which also yielded higher accuracy for longer passages.

From our experiments with text-matching software, we conclude that if plagiarists copy a few paragraphs and employ a paraphrasing tool to obfuscate their misconduct, the similarity is often below the text-matching tool’s reporting threshold, thus causing the plagiarism to remain undetected. Our classification approaches for machine-paraphrased text can be a valuable complement

Table 8. An illustrative sample of three examples for each paraphrasing tool, data source, and classification method.

Original Paragraphs:

- A mathematically rigorous approach to quantum field theory based on operator algebras is called an algebraic quantum field theory...
- "Nuts" contains 5 instrumental compositions written and produced by Streisand, with the exception of "The Bar", including additional writing from Richard Baskin. All of the songs were recorded throughout 1987...
- Most of activities are carried out internally using internal resources. The cost and financial demandingness of the project are presented in the table below...

SpinBot Paraphrased	Source	Turnitin[†]	PlagScan[†]	ML*	NLM*
A numerically thorough way to deal with quantum field hypothesis dependent on administrator algebras is called an arithmetical quantum field hypothesis...	arXiv	25.30	0.00	73.11	99.99

SpinnerChief-IF Paraphrased	Source	Turnitin[†]	PlagScan[†]	ML*	NLM*
"Nuts" consists of five instrumental compositions written and created by Streisand, with the exception of "The Bar", which includes extra writing from Rich Baskin. All of the music were recorded in 1987...	Wiki	35.80	23.90	98.87	72.46

SpinnerChief-DF Paraphrased	Source	Turnitin[†]	PlagScan[†]	ML*	NLM*
The majority of activities are carried out internally making use of internal resources. The cost and economic demandingness of the project are illustrated in the table below...	Thesis	0.00	0.00	49.81	63.81

- > [†]text-match in %.
- > *prediction score in % for the best performing models Longformer and w2v+NB (see Section 4).
- > Red background highlights changed tokens of the original version.
- > Ellipsis ("...") indicates the remainder of the paragraph.

to text-matching software. The additional analysis step could alert users when indicators of machine-obfuscated text have been identified.

We provide an illustrative example for text from arXiv, Wikipedia, and theses, their modified versions using SpinBot and SpinnerChief, and classification scores using text-matching software (Turnitin, PlagScan), the best performing neural language model (Longformer), and the best combination of machine learning classifier and word embeddings (w2v+NB) in Table 8.

5 Conclusion

In this paper, we analyze two new collections (arXiv and theses), an additional paraphrasing tool (SpinnerChief), eight neural language models based on the Transformer architecture (Table 6), and two popular text-matching systems (Turnitin and PlagScan). We selected training and test sets that reflect documents particularly relevant for the plagiarism detection use case. The arXiv collection represents scientific papers written by expert researchers. Graduation theses of non-native English speakers provide writing samples of authors whose style varies considerably. Wikipedia articles represent collaboratively authored documents for many topics and one of the sources from which students plagiarize most frequently.

We investigated the use of traditional, pre-trained word embedding models in combination with machine learning classifiers and recent neural language models. For eight of our nine test sets, Transformer-based techniques proposing changes in the training architecture achieved the highest scores. In particular, Longformer achieved the best classification performance overall.

Transferring the classifiers trained on the SpinBot training set to SpinnerChief test sets caused a drop in the approaches' average classification performance. For SpinnerChief-IF, a test set exhibiting a similar ratio of altered words as the training set, the average F1-Micro scores of the best-performing classifiers dropped by approx. 21%, that of human evaluators by approx. 13%. However, the best-performing models were still capable of classifying machine-paraphrased paragraphs with F1-Micro scores ranging between 74.8% to 80.5%. We partially attribute the loss in performance in recognizing SpinnerChief test cases to the obfuscation's strength and not exclusively to deficiencies in the transferred classifiers.

We showed that our approaches can complement text-matching software, such as PlagScan and Turnitin, which often fail to identify machine-paraphrased plagiarism. The main advantage of machine learning models over text matching software is the models' ability to identify machine-paraphrased text even if the source document is not accessible to the detection system. The classification approaches we investigated could be integrated as an additional step within the detection process of text-matching software to alert users of likely machine-paraphrased text. The presence of such obfuscated text is a strong indicator of deliberate misconduct.

The classification approaches we devised are robust to identifying machine-paraphrased text, which educators face regularly. To support practitioners and facilitate an extension of the research on this important task, the data²² and code²³ of our study, as well as a web-based demonstration system for the best-performing machine learning classifier²⁴ (NB+w2v) and neural language model²⁵ (Longformer) are openly available.

6 Future Work

Our experiments indicate that obtaining additional training data is a promising approach for improving artificial intelligence-backed approaches for identifying machine-paraphrased text. Additional training data should cover more paraphrasing tools, topics, and languages. We see a community-driven open data effort as a promising option for generating a comprehensive training set. We encourage researchers investigating machine-paraphrase detection to share their data and contribute to the consolidation and extension of datasets, such as the one we publish with this paper.

²²<https://doi.org/10.5281/zenodo.3608000>

²³<https://github.com/jpelhaW/ParaphraseDetection>

²⁴<http://purl.org/spindetector>

²⁵<https://huggingface.co/jpelhaw/longformer-base-plagiarism-detection>

Obtaining effective training data is challenging due to many paraphrasing tools, nontransparent paraphrasing approaches, frequent interconnections of paraphrasing tools, and the questionable business model of the tool providers. If paying paraphrasing services to obtain data proves prohibitive, a crowdsourcing effort could overcome the problem. Another interesting direction would be to use auto-encoding language models to paraphrase text or generate new text with auto-regressive models. This setup will be more realistic in the future as language models are publicly available and generate text that is difficult to distinguish from human writing.

References

1. Alvi, F., Stevenson, M., Clough, P.: Paraphrase type identification for plagiarism detection using contexts and word embeddings. *International Journal of Educational Technology in Higher Education* **18**(1), 42 (Aug 2021). <https://doi.org/10.1186/s41239-021-00277-8>, <https://doi.org/10.1186/s41239-021-00277-8>
2. Beltagy, I., Lo, K., Cohan, A.: SciBERT: A Pretrained Language Model for Scientific Text. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. pp. 3613–3618. Association for Computational Linguistics, Hong Kong, China (2019). <https://doi.org/10/ggcgtm>
3. Beltagy, I., Peters, M.E., Cohan, A.: Longformer: The Long-Document Transformer. *arXiv:2004.05150 [cs]* (Apr 2020)
4. Bojanowski, P., Grave, E., Joulin, A., Mikolov, T.: Enriching Word Vectors with Subword Information. *Transactions of the Association for Computational Linguistics* **5**, 135–146 (Dec 2017). <https://doi.org/10/gfw9cs>
5. Clark, K., Luong, M.T., Le, Q.V., Manning, C.D.: ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators. *arXiv:2003.10555 [cs]* (Mar 2020)
6. Conneau, A., Kiela, D., Schwenk, H., Barrault, L., Bordes, A.: Supervised learning of universal sentence representations from natural language inference data. *Proceedings Conference on Empirical Methods in Natural Language Processing* (2017). <https://doi.org/10.18653/v1/d17-1070>
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv abs/1810.04805* (2018)
8. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv:1810.04805 [cs]* (May 2019)
9. Dey, K., Shrivastava, R., Kaushik, S.: A Paraphrase and Semantic Similarity Detection System for User Generated Short-Text Content on Microblogs. In: *Proceedings International Conference on Computational Linguistics (Coling)*. pp. 2880–2890 (2016)
10. Dolan, W.B., Brockett, C.: Automatically constructing a corpus of sentential paraphrases. In: *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)* (2005)
11. Foltýnek, T., Dlabolová, D., Anohina-Naumeca, A., Razi, S., Kravjar, J., Kamzola, L., Guerrero-Dib, J., Çelik, Ö., Weber-Wulff, D.: Testing of support tools for

- plagiarism detection. *International Journal of Educational Technology in Higher Education* **17**(1), 46 (2020). <https://doi.org/10.1186/s41239-020-00192-4>
12. Foltýnek, T., Meuschke, N., Gipp, B.: Academic Plagiarism Detection: A Systematic Literature Review. *ACM Computing Surveys* **52**(6), 112:1–112:42 (2019). <https://doi.org/10.1145/3345317>
 13. Foltýnek, T., Ruas, T., Scharpf, P., Meuschke, N., Schubotz, M., Grosky, W., Gipp, B.: Detecting Machine-obfuscated Plagiarism. In: *Proceedings of the iConference 2020*. LNCS, Springer (2020)
 14. Gharavi, E., Veisi, H., Rosso, P.: Scalable and language-independent embedding-based approach for plagiarism detection considering obfuscation type: no training phase. *Neural Computing and Applications* **32**(14), 10593–10607 (2020). <https://doi.org/10.1007/s00521-019-04594-y>
 15. Gutmann, M., Hyvärinen, A.: Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In: *Proc. Int. Conf. on Artificial Intelligence and Statistics (AISTATS)*. JMLR W&CP, vol. 9, pp. 297–304 (2010)
 16. Hunt, E., Janamsetty, R., Kinares, C., Koh, C., Sanchez, A., Zhan, F., Ozdemir, M., Waseem, S., Yolcu, O., Dahal, B., Zhan, J., Gewali, L., Oh, P.: Machine learning models for paraphrase identification and its applications on plagiarism detection. *Proceedings 10th IEEE International Conference on Big Knowledge* pp. 97–104 (2019). <https://doi.org/10.1109/ICBK.2019.00021>
 17. Iyer, S., Dandekar, N., Csernai, K.: First quora dataset release: Question pairs (2017), <https://data.quora.com/First-Quora-Dataset-Release-Question-Pairs>
 18. Lan, W., Qiu, S., He, H., Xu, W.: A Continuously Growing Dataset of Sentential Paraphrases. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. pp. 1224–1234. Association for Computational Linguistics, Copenhagen, Denmark (2017). <https://doi.org/10.18653/v1/D17-1126>
 19. Lan, W., Xu, W.: Neural Network Models for Paraphrase Identification, Semantic Textual Similarity, Natural Language Inference, and Question Answering. arXiv:1806.04330 [cs] (Aug 2018)
 20. Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., Soricut, R.: ALBERT: A Lite BERT for Self-supervised Learning of Language Representations. arXiv:1909.11942 [cs] (Sep 2019)
 21. Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., Soricut, R.: Albert: A lite bert for self-supervised learning of language representations. arXiv:1909.11942 [cs] (2019)
 22. Lau, J.H., Baldwin, T.: An empirical evaluation of doc2vec with practical insights into document embedding generation. In: *Proceedings Workshop on Representation Learning for NLP* (2016). <https://doi.org/10.18653/v1/w16-1609>
 23. Le, Q., Mikolov, T.: Distributed representations of sentences and documents. In: *Proceedings 31st International Conference on Machine Learning*. vol. 32, pp. 1188–1196 (2014)
 24. Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., Zettlemoyer, L.: BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. arXiv:1910.13461 [cs] (Oct 2019)
 25. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: RoBERTa: A Robustly Optimized BERT Pretraining Approach. arXiv:1907.11692 [cs] (Jul 2019)

26. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv:1907.11692 [cs] (2019)
27. Meuschke, N.: Analyzing Non-Textual Content Elements to Detect Academic Plagiarism. Doctoral Thesis, University of Konstanz, Dept. of Computer and Information Science, Konstanz, Germany (2021). <https://doi.org/10.5281/zenodo.4913345>
28. Meuschke, N., Gondek, C., Seebacher, D., Breitingner, C., Keim, D., Gipp, B.: An Adaptive Image-Based Plagiarism Detection Approach. In: Proceedings 18th ACM/IEEE Joint Conference on Digital Libraries. pp. 131–140 (2018). <https://doi.org/10.1145/3197026.3197042>
29. Meuschke, N., Stange, V., Schubotz, M., Gipp, B.: HyPlag: A Hybrid Approach to Academic Plagiarism Detection. In: Proceedings 41st International ACM SIGIR Conference on Research & Development in Information Retrieval. pp. 1321–1324 (2018). <https://doi.org/10.1145/3209978.3210177>
30. Meuschke, N., Stange, V., Schubotz, M., Kramer, M., Gipp, B.: Improving Academic Plagiarism Detection for STEM Documents by Analyzing Mathematical Content and Citations. In: Proceedings ACM/IEEE Joint Conference on Digital Libraries. pp. 120–129 (2019). <https://doi.org/10.1109/JCDL.2019.00026>
31. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., Dean, J.: Distributed Representations of Words and Phrases and their Compositionality. arXiv:1310.4546 [cs, stat] (Oct 2013)
32. Napoles, C., Gormley, M., Van Durme, B.: Annotated Gigaword. In: Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction (AKBC-WEKEX). pp. 95–100. Association for Computational Linguistics, Montréal, Canada (Jun 2012)
33. Ostendorff, M., Ash, E., Ruas, T., Gipp, B., Moreno-Schneider, J., Rehm, G.: Evaluating document representations for content-based legal literature recommendations. In: Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law. pp. 109–118. ACM, São Paulo Brazil (Jun 2021). <https://doi.org/10.1145/3462757.3466073>, <https://arxiv.org/pdf/2104.13841.pdf>
34. Ostendorff, M., Ruas, T., Blume, T., Gipp, B., Rehm, G.: Aspect-based Document Similarity for Research Papers. In: Proceedings of the 28th International Conference on Computational Linguistics. pp. 6194–6206. International Committee on Computational Linguistics, Barcelona, Spain (Online) (2020). <https://doi.org/10.18653/v1/2020.coling-main.545>, <https://aclanthology.org/2020.coling-main.545.pdf>
35. Pennington, J., Socher, R., Manning, C.D.: Glove: Global vectors for word representation. In: Proceedings Conference on Empirical Methods in Natural Language Processing. vol. 14, pp. 1532–1543 (2014). <https://doi.org/10/gfshwg>
36. Perone, C.S., Silveira, R., Paula, T.S.: Evaluation of sentence embeddings in downstream and linguistic probing tasks. arXiv **abs/1806.06259** (2018)
37. Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., Zettlemoyer, L.: Deep Contextualized Word Representations. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers). pp. 2227–2237. Association for Computational Linguistics, New Orleans, Louisiana (2018). <https://doi.org/10.18653/v1/n18-1202>
38. Prentice, F.M., Kinden, C.E.: Paraphrasing Tools, Language Translation Tools and Plagiarism: An Exploratory Study. *International Journal for Educational Integrity* **14**(1) (2018). <https://doi.org/10.1007/s40979-018-0036-7>

39. Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I.: Language models are unsupervised multitask learners (2019)
40. Rogerson, A.M., McCarthy, G.: Using Internet based paraphrasing tools: Original work, patchwriting or facilitated plagiarism? *International Journal for Educational Integrity* **13**(1) (2017). <https://doi.org/10.1007/s40979-016-0013-y>
41. Ruas, T., Ferreira, C.H.P., Grosky, W., de França, F.O., de Medeiros, D.M.R.: Enhanced word embeddings using multi-semantic representation through lexical chains. *Information Sciences* **532**, 16–32 (Sep 2020). <https://doi.org/10.1016/j.ins.2020.04.048>
42. Ruas, T., Grosky, W., Aizawa, A.: Multi-sense embeddings through a word sense disambiguation process. *Expert Systems with Applications* **136**, 288–303 (Dec 2019). <https://doi.org/10.1016/j.eswa.2019.06.026>
43. Ruas, T., Grosky, W., Aizawa, A.: Multi-sense embeddings through a word sense disambiguation process. *Expert Systems with Applications* **136**, 288 – 303 (2019). <https://doi.org/https://doi.org/10.1016/j.eswa.2019.06.026>
44. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter. arXiv:1910.01108 [cs] (Oct 2019)
45. Spinde, T., Plank, M., Krieger, J.D., Ruas, T., Gipp, B., Aizawa, A.: Neural Media Bias Detection Using Distant Supervision With BABE - Bias Annotations By Experts. In: Findings of the Association for Computational Linguistics: EMNLP 2021. Dominican Republic (Nov 2021), tex.pubstate: published tex.tppubtype: in-proceedings
46. Subramanian, S., Trischler, A., Bengio, Y., Pal, C.J.: Learning General Purpose Distributed Sentence Representations via Large Scale Multi-task Learning. arXiv:1804.00079 [cs] (Mar 2018)
47. Trinh, T.H., Le, Q.V.: A simple method for commonsense reasoning. arXiv:1806.02847 [cs] (2019)
48. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: Attention is all you need. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 30, pp. 5998–6008. Curran Associates, Inc. (2017), <https://arxiv.org/abs/1706.03762>
49. Wahle, J.P., Ashok, N., Ruas, T., Meuschke, N., Ghosal, T., Gipp, B.: Testing the Generalization of Neural Language Models for COVID-19 Misinformation Detection. In: *Proceedings of the iConference* (February 2022)
50. Wahle, J.P., Ruas, T., Meuschke, N., Gipp, B.: Are Neural Language Models Good Plagiarists? A Benchmark for Neural Paraphrase Detection. In: *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries (JCDL)*. IEEE, Washington, USA (Sep 2021)
51. Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., Bowman, S.R.: GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. arXiv:1804.07461 [cs] (Feb 2019)
52. Weber-Wulff, D.: Plagiarism detectors are a crutch, and a problem. *Nature* (2019). <https://doi.org/10.1038/d41586-019-00893-5>
53. Xu, W.: Data-Drive Approaches for Paraphrasing Across Language Variations. Ph.D. thesis, Department of Computer Science, New York University (2014), <http://www.cis.upenn.edu/~xwe/files/thesis-wei.pdf>
54. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R., Le, Q.V.: XLNet: Generalized Autoregressive Pretraining for Language Understanding. arXiv:1906.08237 [cs] (Jun 2019)

55. Zellers, R., Holtzman, A., Rashkin, H., Bisk, Y., Farhadi, A., Roesner, F., Choi, Y.: Defending against neural fake news. arXiv:1905.12616 [cs] (2019)
56. Zhang, Q., Wang, D.Y., Voelker, G.M.: Dspin: Detecting automatically spun content on the web. In: Proceedings Network and Distributed System Security (NDSS) Symposium. pp. 23–26 (2014). <https://doi.org/10.14722/ndss.2014.23004>
57. Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urtasun, R., Torralba, A., Fidler, S.: Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In: The IEEE International Conference on Computer Vision (ICCV) (December 2015)