

# Investigating the detection of Tortured Phrases in Scientific Literature

Puthineath Lay<sup>1</sup>, Martin Lentschat<sup>1</sup>, and Cyril Labbé<sup>1</sup>

<sup>1</sup> Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, 38000 Grenoble, France  
[puthineath.lay@cadt.edu.kh](mailto:puthineath.lay@cadt.edu.kh), [martin.lentschat@univ-grenoble-alpes.fr](mailto:martin.lentschat@univ-grenoble-alpes.fr)

## Abstract

With the help of online tools, unscrupulous authors can today generate a pseudo-scientific article and attempt to publish it. Some of these tools work by replacing or paraphrasing existing texts to produce new content, but they have a tendency to generate nonsensical expressions. A recent study introduced the concept of “tortured phrase”, an unexpected odd phrase that appears instead of the fixed expression. E.g. *counterfeit consciousness* instead of *artificial intelligence*. The present study aims at investigating how tortured phrases, that are not yet listed, can be detected automatically. We conducted several experiments, including non-neural binary classification, neural binary classification and cosine similarity comparison of the phrase tokens, yielding noticeable results.

## 1 Introduction

Scientific texts generated by computer programs can be meaningless, and fake generated papers are served and sold by various publishers with the estimation of 4.29 documents every one million reports (Cabanac and Labbé, 2021). But generated texts are also meaningful: with the inputs of a thousand articles, new books are now produced (e.g. Beta Writer, 2019). Despite the ability of text-generators to produce counterfeit publications, meaningless generated papers can be easily spotted by both machines and humans (Cabanac et al., 2021). Texts produced by neural language models are more difficult to spot (Hutson et al., 2021). These neural language models can produce paraphrased texts that are closer to human-written texts (Brown et al., 2020), and therefore machine-paraphrased texts are harder to differentiate from the human-written texts.

Online tools such as Spinbot, and Spinner Chief are used to paraphrase texts. However the capacity of a paraphrasing software to assist a writer can be harmful to the scientific literature. Cabanac et al.

(2021) screened recent publications (e.g. in the journal *Microprocessors and Microsystems*) and discovered over 500 meaning less phrases in those scientific papers. They called it “tortured phrases”, unexpected odd phrases replacing the lexicalised expression, such as *counterfeit consciousness* instead of *artificial intelligence* (i.e., the expected phrase). The database of tortured phrases, and articles that contain them, have since been expanded to over 9000 publications in different domains such as Computer Sciences, Biology or Medicine.

In this paper, we investigate strategies to automatically detect new (i.e. unlisted) tortured phrases. Focusing solely on tortured phrases detection, and not paraphrased text in general, we will use recent machine learning techniques and state-of-the-art language models. Our methods were trained on a corpus composed of 141 known tortured phrases, taking their sentences as contexts, and aims at detecting never-seen-before tortured phrases. All code and corpus used are available online.

## 2 Related Works

Up to now, no dataset has been built for the automatic detection of tortured phrases. In Cabanac et al. (2021), authors and contributors collected a set of tortured phrases and their expected phrases that we will use as dataset. Wahle et al. (2021) used Spinbot and Spinnerchief to paraphrase original data from several sources such as an arXiv test sets, graduation theses, and Wikipedia articles. Their study aims at detecting whether a paragraph is machine-paraphrased or not. The authors tested classic machine learning approaches and neural language models based on the Transformer architecture (Vaswani et al., 2017), such as BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019), ALBERT (Lan et al., 2019), Longformer (Beltagy et al., 2020), and others. They showed that such approaches can complement text-matching software, such as PlagScan and Turnitin, which often fail to

notice machine-paraphrased plagiarism.

Because paraphrasing tools like Spinbot and Spinnerchief can generate tortured phrases, the dataset created by Wahle et al. (2021) surely contains such phrases. But the task we aim at, i.e. detecting new tortured phrases, is more specific than detecting paraphrased text. Thus, we investigated three supervised machine learning classifiers: Random Forest, Perceptron, and Transformer-based model. Random Forest classifier (Breiman, 2001) is an ensemble learning method that builds decision trees and classifies each data according to the most selected class. Perceptron (Rosenblatt, 1958) is a linear classifier used to classify vectors of numbers. Term Frequency-Inverse Document Frequency (TF-IDF), GloVe (Pennington et al., 2014), and BERT (Devlin et al., 2018) were used for word vector representation. These models were chosen for their state-of-the-art performances and to compare how a model with fixed vectors (i.e. GloVe) compares with a model using dynamic ones (i.e. BERT).

### 3 Building datasets of tortured phrases

This experiment uses two data sources: the tortured phrases (Cabanac et al., 2021) database and the contexts containing tortured phrases from Wahle et al. (2021). We automatically extracted the context of known tortured phrases from the corpus of Wahle et al. (2021) to build a training set.

Tortured phrases identified by Cabanac et al. (2021) consists of 558 tortured phrases (e.g. Table 1). These phrases were annotated by the authors and other contributors in several media (e.g. PubPeer, Twitter) in 2021-2022. After pre-processing, we retained 545 tortured phrases.

Tortured phrases	Expected phrases
innocent Bayes	naive Bayes
ghostly grouping	spectral clustering
Unused Britain	New England
Joined together states	United States
immature nations	developing countries

Table 1: Example of tortured and expected phrases.

The dataset of Wahle et al. (2021) is constituted of 193,646 paragraphs, paraphrased using Spinbot and Spinnerchief. 65,433 original data were retrieved from several sources: arXiv, graduation theses of ESL students at the Mendel University in Brno (Czech Republic) and Wikipedia articles.

We extracted the paragraphs containing tortured phrases (Cabanac et al., 2021) to build a training and evaluation corpus. This resulted in 1,104 paragraphs with tortured phrases and 1,668 paragraphs without tortured phrases (randomly extracted from the non-paraphrased original data).

**Data augmentation: five-grams extraction** To increase the training data, we extracted the n-grams (i.e. sequences of n adjacent tokens) of each sentence, with  $n = 5$ , as it is the maximum length of the known tortured phrases. A five-gram is considered positive if a complete tortured phrase appears in that five-gram. This produced 38,397 five-grams, 5,024 positive five-grams (in the '1' class) and 33,373 negative five-grams (in the '0' class).

## 4 Experiment and Result

We investigated binary classifiers to check the difference between paragraphs or five-grams containing tortured phrases in several settings. The paragraphs and five-grams containing tortured phrases are considered positives, with label '1', while negative paragraphs are labeled as '0'. Accuracy, precision, recall, and F-measure are used to evaluate the classification performances.

**Classifiers: Random Forest and Perceptron** In this experiment, five-grams data are used in the classification. The five-grams are converted to a numerical representation using Sklearn TF-IDF count vectorizer and split randomly 80% for training and 20% for testing. We used the Scikit Learn library for the Random Forest and Perceptron with the default value of all parameters.

The result in Table 2 shows an accuracy for the Random Forest classifier of .98 and the Perceptron of .94. The precision, recall, and F1-score of the Random Forest classifier are high, especially in class 0, and the results in class 1 is slightly lower than in class 0. The precision, recall, and F1-score of the Perceptron method is slightly lower than that of the Random Forest classifier, but it is still comparable for class 0. In class 1, Perceptron results are significantly lower compared to the results of Random Forest classifier. We also observe results higher in class 0 than in class 1, this might be due to the data imbalance.

After observing the accuracy, precision, recall, and F1-score, we see that the models perform well based on TF-IDF vector representation. However, it is believable that the models learned to classify

five-grams based solely on specific words: since the training and test data were split randomly, a tortured phrase can be present in both sub-sets.

#### **Transformer-based classifier on paragraphs**

Here, we seek at detecting paragraphs containing at least one tortured phrase. The data are split 67% for training and 33% for test set. The architecture of this model is based on the Transformer technique. Pre-trained transformers from Huggingface, `distilbert-base-uncased` model (Sanh et al., 2019), was chosen for its lightness and speed. We applied transfer learning by adding one linear layer for classification purposes. In that linear layer, the number of input features was set to 768 with an output size of 2, indicating class 0 and class 1. The model was trained on 10 epochs.

The results in Table 2 show an accuracy of .86. The .92 precision on class 1 is higher than on class 0. For the recall and F1-score, class 0 gets a better result than class 1. Since the amount of paragraph data are small, we suspect that only a few tokens constitute the tortured phrases, and that the rest of the token’s paragraph affect the performance.

#### **Transformer-based classifier on five-grams**

The data of class ‘1’ was split 79% for training and 21% for testing by filtering the test set with tortured phrases not present in the training set. In this experiment, two versions of the model were trained: one using the entire dataset and one with a proportion of data balanced in both classes.

The training data are made of 28,995 five-grams (25,029 in class ‘0’ and 3,966 in class ‘1’). The test data are made of 9,402 five-grams (8,344 and 1,058). Table 2 shows an accuracy of .88. Precision, recall, and F1-score on class ‘1’ are exceptionally low compared to the ‘0’ class.

Regarding the classifier with balanced data, the size of the training set is 7,932 five-grams (3,966 in each class) and the size of the testing set is 2,116 five-grams (1,058 in each class). The accuracy is .71, and the precision, recall, and F1-score are around .70 in both classes (cf. bold values in Table 2). In this experiment, the balance of the classes induces a greater reliability of the results and we believe this approach presents the best applicability. The model focuses on the tortured phrases in five-grams rather than tortured phrases in the whole paragraph, so the model can learn to generalize the five-grams containing tortured phrases or not.

## **4.1 Cosine similarity comparison**

We studied the cosine similarity between tokens in the tortured phrases compared to the similarity of tokens in the expected phrases. Cabanac et al. (2021) annotated a dataset of tortured phrases, and their respective expected phrases, from several media such as PubPeer during 2021-2022.

We intuitively expect that the cosine score of the tokens in the phrases could yield noticeable results, useful to differentiate tortured and expected phrases due to the similarity, or non-similarity, of adjacent tokens. The expected phrases are idioms (i.e. multi-word-expression forming a lexical and semantic unit) and, as such, we hypothesize that the semantic score defined by the cosine of the vectors between their terms should be higher than for the tortured phrases, which words are less likely to be semantically related or frequently associated. If validated, such observation could help distinguish tortured phrases from legitimate ones. For this experiment, only the similarity of two-tokens phrases were computed, using two kinds of word embedding models: GloVe and BERT.

**Cosine similarity on phrases using BERT** In this study, we used BERT (Devlin et al., 2018) as the word embedding model. We followed the architecture of McCormick (2019) by summing the last four layers of 12 layers of BERT to get one-word vectors with 768 values.

Since the BertTokenizer will separate unknown words into sub-words (e.g. *vitality utilize* becomes *vital*, *#ity*, and *utilize*), it can be complicated to compute their cosine similarity. We chose to discard tortured phrases containing words unknown by the model. We retained 82 tortured phrases after the tokenization process.

The scores obtained from cosine computation between word pairs in tortured and expected phrases present slight differences, as shown in Figure 1. The median scores of expected and tortured phrases are .51 and .49, respectively. The absence of significant differences can be explained by the nature of the BERT model: a two-word context is probably not sufficient to differentiate tortured and expected phrases using cosine similarity.

#### **Cosine similarity on phrases using GloVe**

For this experiment, we computed the cosine similarity of token pairs in tortured and expected phrases. We used the pre-trained GloVe word embedding (Pennington et al., 2014)

Classifiers class	Data type	Accuracy	Precision		Recall		F1-score	
			0	1	0	1	0	1
Random Forest	Random five-grams	.98	.99	.92	.99	.91	.99	.92
Perceptron	Random five-grams	.94	.96	.84	.98	.69	.97	.75
Transformer	Paragraph	.86	.82	.92	.94	.77	.87	.84
Transformer	Random five-grams	.88	.89	.42	.99	.03	.93	.06
Transformer	Balanced five-grams	<b>.71</b>	<b>.67</b>	<b>.75</b>	<b>.79</b>	<b>.62</b>	<b>.73</b>	<b>.68</b>

Table 2: Classification results.

(glove.6B.200d.txt) which was trained on Wikipedia 2014 and Gigaword. Unlike BERT, GloVe is a context-free model, meaning that each word in this pre-trained model is assigned to one constant vector. However, GloVe vocabulary is limited and thus some tokens in the phrases might not appear in this model. For this issue, we padded the out-of-vocabulary word with 0. For the phrases in which both words do not appear in the model, the cosine similarity is 0. We discarded the phrases whose scores were lower than or equal to 0 ( $\text{cosine\_score} \leq 0$ ). As a result, 139 phrases are used for this experiment.

The Figure 1 indicates that the cosine similarity scores of tortured phrases tend to be smaller than those of expected phrases when using GloVe. The median score of expected phrases is .3 and the median score of tortured phrases is .12.




Figure 1: Comparison of cosine score of phrases using BERT and GloVe.

These results indicate that the cosine score between terms could be employed to differentiate tortured phrases from legitimate ones. Since Bert relies on embeddings, it is overly influenced by the phrase contexts to yields useful results. With Glove, or another language model with static vectors, one could chose a threshold to classify phrases, e.g. as legitimate, tortured, or requiring human expertise.

## 5 Conclusions

In this research, we aimed at detecting new tortured phrases. We studied different classification approaches and examined the characteristics of tortured phrases using cosine similarity. The result of Perceptron and Random Forest classifier are high, but we intuitively suspect they are not reliable due to the word representation using TF-IDF vectorization. The Transformer-based classifier model with paragraph data provided the best result among Transformer models. However, we suspect that the model learned to classify paragraphs based only on a few tokens, and classifying paragraphs is not sufficient to detect the exact tortured phrases. Thus, the Transformer-based classifier model five-gram data yields the best result with balanced classes (i.e. results above .70 for all metrics).

We also studied the use of cosine similarity between the phrase tokens to identify new tortured phrase. This showed that language model with fixed vectors (e.g. Glove) could be used to classify part of the phrases.

Future research should include more human evaluation of tortured phrases and a bigger dataset tortured phrases with their context. To improve the classification, future work could investigate Support Vector Machine (SVM) and Naïve Bayes model (NB): SVM performs better with a small dataset and binary class, while NB can provide probabilities of a prediction. Finally, computing the cosine similarity of the tortured phrase and expected phrase pairs within the whole context to see tortured phrases' performance.

## Acknowledgments

The NanoBubbles project has received Synergy grant funding from the European Research Council (ERC), within the European Union's Horizon 2020 program, grant agreement no. 951393.

## References

- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Lithium-Ion Beta Writer. 2019. *a Machine-Generated Summary of Current Research/Beta Writer*. Springer International Publishing.
- Leo Breiman. 2001. Random forests. *Machine learning*, 45(1):5–32.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Guillaume Cabanac and Cyril Labbé. 2021. Prevalence of nonsensical algorithmically generated papers in the scientific literature. *Journal of the Association for Information Science and Technology*, 72(12):1461–1476.
- Guillaume Cabanac, Cyril Labbé, and Alexander Magazinov. 2021. Tortured phrases: A dubious writing style emerging in science. evidence of critical issues affecting established journals. *CoRR*, abs/2107.06751.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Matthew Hutson et al. 2021. Robo-writers: the rise and risks of language-generating ai. *Nature*, 591(7848):22–25.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Chris McCormick. 2019. [Bert word embeddings tutorial](#).
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. [Glove: Global vectors for word representation](#). In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.
- Frank Rosenblatt. 1958. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Jan Wahle, Terry Ruas, Tomas Foltyněk, Norman Meuschke, and Bela Gipp. 2021. Identifying machine-paraphrased plagiarism.