

# DeepPavlov at SemEval-2024 Task 8: Leveraging Transfer Learning for Detecting Boundaries of Machine-Generated Texts

Anastasia Voznyuk and Vasily Konovalov

Moscow Institute of Physics and Technology

{vozniuk.ae, vasily.konovalov}@phystech.edu

## Abstract

The Multigenerator, Multidomain, and Multilingual Black-Box Machine-Generated Text Detection shared task in the SemEval-2024 competition aims to tackle the problem of misusing collaborative human-AI writing. Although there are a lot of existing detectors of AI content, they are often designed to give a binary answer and thus may not be suitable for more nuanced problem of finding the boundaries between human-written and machine-generated texts, while hybrid human-AI writing becomes more and more popular. In this paper, we address the boundary detection problem. Particularly, we present a pipeline for augmenting data for supervised fine-tuning of DeBERTaV3. We receive new best MAE score, according to the leaderboard of the competition, with this pipeline.

## 1 Introduction

Recently, there has been a rapid development of auto-regressive language models, for example, GPT-3 (Brown et al., 2020), GPT-4 (OpenAI, 2023), and LLaMA2 (Touvron et al., 2023). These models are trained on enormous amounts of data and are able to produce coherent texts that can be indistinguishable from human-written texts (Dugan et al., 2022).

The SemEval-2024 Task 8 competition suggests to tackle the problem of detecting machine-generated texts. This problem has become more relevant recently due to the release of ChatGPT<sup>1</sup>, a model by OpenAI that simplified the access to the large language models (LLM) and their usage. For example, LLM can be maliciously used to generate fake news (Zellers et al., 2019). There are also some concerns raised among scientists (Ma et al., 2023) and educators (Zeng et al., 2023) that the usage of LLMs will devalue the process of learning and research.

<sup>1</sup><https://openai.com/blog/chatgpt>

The commonly used approach to formulate the task of detecting machine-generated texts is a binary classification task (Jawahar et al., 2020). In this case, a text can be attributed to either a human or a LLM. Otherwise, the task can be formulated as a multiclass classification or an authorship attribution task (Uchendu et al., 2020), where it is needed to determine which one of the  $k$  authors is the real author of the given text. Finally, the trend toward human-AI collaborative writing is rising, which highlights the importance of the boundary detection task. In this setup, text contains consecutive chunks of different authorship, and it is required to detect where the boundaries between chunks lie and who is the author of every chunk. Due to its complexity, it is usually assumed the text has a human-written prefix and the rest of the text is AI-generated (Dugan et al., 2022; Cutler et al., 2021; Kushnareva et al., 2023).

Our main contributions are three-fold:

1. We receive the new best MAE score on the task of detecting the boundary between human-written and machine-generated parts of the text.
2. We present a new simple yet effective pipeline of augmenting data for the task of boundary detection, which allows us to get more data for training and improve the results of fine-tuning large language models.
3. We compare the performance of several fine-tuned models with different architectures on various amounts of training data.

Additionally, we've made the code of augmentation publicly available.<sup>2</sup>

<sup>2</sup><https://github.com/natriiistorm/semEval2024-boundary-detection>

## 2 Related Work

Most of approaches (Jawahar et al., 2020) for machine-generated text detection are based on calculating linguistic (Fröhling and Zubiaga, 2021), stylometric, and statistical features, as well as on using classical machine learning methods like logistic regression, random forest, and gradient boosting as classifiers. Among commonly used features are word and n-gram frequencies (Manjavacas et al., 2017), and tf-idf (Solaiman et al., 2019).

An alternative strategy is to use zero-shot techniques based on the internal metrics of the texts. For example, token-wise log probability can be evaluated by models like GROVER (Zellers et al., 2019) or GPT-2 (Solaiman et al., 2019). A probability threshold is established to distinguish writings produced by machines from those written by humans. Moreover, rank (Gehrmann et al., 2019) or log-rank (Mitchell et al., 2023) can be calculated for each token and compared for consistency with the prior context.

It's shown by Ippolito et al. (2020) that feature-based methods are inferior to methods based on using encoders of pretrained language models like BERT (Devlin et al., 2019) as a basis for fine-tuning on the selected domain. Representations from auto-regressive language models can be used as input for the classification head. Such transformer-based methods require supervised detection examples for further training. Among the models commonly used for fine-tuning are RoBERTa (Liu et al., 2019) and XLNet (Yang et al., 2019). Fine-tuned RoBERTa has shown the state-of-the-art results on the problems of authorship attribution (Uchendu et al., 2021).

To tackle the problem of mixed human-machine writing, Dugan et al. (2022) introduces the Real Or Fake Text (RoFT) tool, where humans were asked to detect the sentence where the text transitions from human-written text to machine-generated text. One of the possible formulations of this task is multilabel classification (Cutler et al., 2021), where the boundary detector needs to determine the first generated sentence, and the number of this sentence is considered the label of the text. In that work, each sentence is processed separately with shallow classification and regression models based on RoBERTa and SROBERTa (Reimers and Gurevych, 2019). That solution perform well in an in-domain setup, but is limited in an out-of-domain setup.

Any solution for tasks about detecting machine-

generated texts should be robust to domain change. The organisers of the SemEval-2024 Task 8 competition claimed to have added new domains to the test set for testing the robustness of participants' solutions. There are several works on performance of detecting methods on out-of-domain setup, such as Kushnareva et al. (2023) and Zeng et al. (2023). Kushnareva et al. (2023) conclude that perplexity-based and topological features appear to be helpful in case of domain shift.

## 3 Data and Task Description

### 3.1 Task Description

The Multigenerator, Multidomain, and Multilingual Black-Box Machine-Generated Text Detection (Wang et al., 2024) is focused on challenging detectors of machine-generated texts. The dataset, provided by organisers, consists of 3 parts:

1. texts of different authorship for subtask A and subtask B;
2. texts with collaborative human-AI writing for subtask C.

This paper focuses only on subtask C, suggesting a solution to differentiate a human-written prefix from the rest of the AI-generated text. The texts for this subtask are generated in the following way: the language model should continue the human-written text, which is given as a prompt. Several examples of texts are presented in Appendix A. The designated evaluation metric for this task is Mean Absolute Error (MAE), which quantifies the absolute distance between the predicted word and the actual word where the switch of authorship between human and LLM occurs.

### 3.2 Data Description

The dataset for this task is derived from the M4 (Wang et al., 2023) dataset, which contains texts of various domains, various languages, and generators. The authors show that current detectors tend to misclassify machine-generated texts as human-written if they're given a text from a different domain. The texts in the train and validation datasets are generated from scientific paper reviews from PeerRead (Kang et al., 2018). The test set partially consists of texts generated from PeerRead. In order to check the robustness of the solutions to domain shift, texts from Outfox, a dataset of LLM-generated student essays (Koike et al., 2023), are

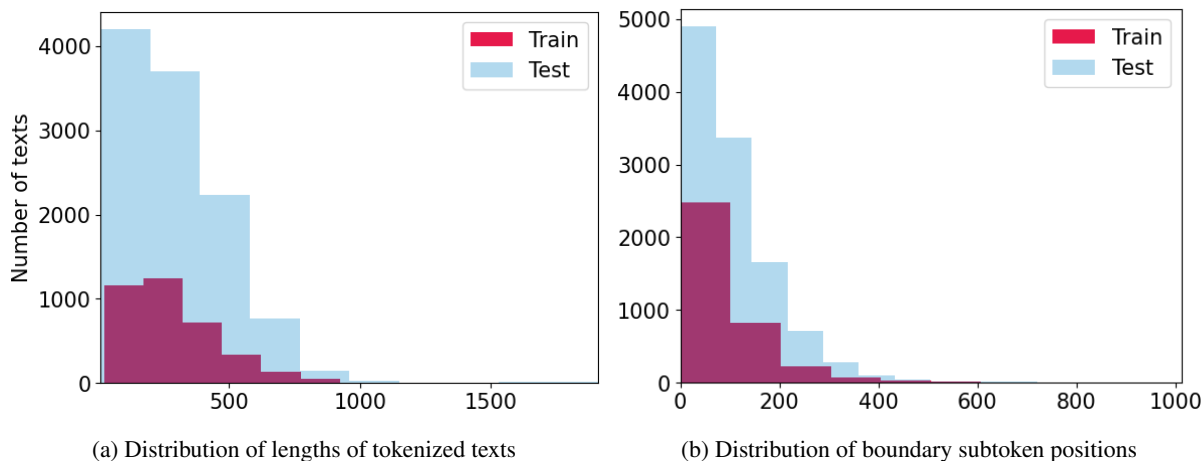


Figure 1: Statistics of the texts in the datasets

added to the dataset. All data sets contain English texts only.

### 3.3 Data Analysis

The distribution of lengths of texts in train and test data sets, tokenized by DeBERTa-V3-base models, is shown in Figure 1a. The majority of the texts in the test set are shorter in length, but there are several texts that exceed 1,000 subtokens. The distribution of positions of a boundary subtoken in texts is shown in Figure 1b. A boundary subtoken is the first subtoken of a boundary word. In both datasets, there is a distinguishable disproportion in the position of a boundary subtoken, as most of them have the boundary subtoken within the first 200 subtokens. Thus, a model, trained only on this data, will give limited results when encountering longer texts with longer human prefixes in them.

## 4 Proposed Method

The provided train dataset consists only of 3,649 texts. It’s well known that an abundance of in-domain training data is crucial for classifier performance (Kononov et al., 2016). However, during the competition period, it was prohibited to use any external data for training, and thus we were limited to working only with the provided dataset. In this case, getting more training data with some kind of augmentation plays a crucial role. We designed an augmentation pipeline and ran all our experiments on two sets of data: the one provided by the organisers of the competition, described in Section 3, and our augmented data.

### 4.1 Data Augmentation

The general idea of our augmentation pipeline is to split the text into distinct sentences and take several consequent sentences from the text with authorship change. It will make new texts coherent and will not mislead models during training (Ostyakova et al., 2023). In addition, to be useful for training, each sequence should contain a sentence with an authorship change.

Another nuance about the augmentation process is the need to correctly determine the boundary label. The boundary label is calculated as a number of whitespace-split words in the human-written prefix. However, the initial dataset contains texts where a pair or even a sequence of words is split only by line breaks or punctuation symbols. Such a sequence of words should be considered as one word when calculating the boundary label. Thus, we have to pay a lot of attention to whitespace characters during augmentation and do not mistakenly append new whitespace characters between words.

We preprocess each text in the dataset for augmentation in the following way:

1. Split the text into sentences by punctuation symbols.
2. Split the sentences themselves by whitespaces into lists of whitespace-split words.
3. Compare the list of whitespace-split words from previous step with the list of whitespace-split words obtained by splitting the text itself. In case of discrepancy, fix it, depending on its type.

See example of preprocessing for the text from train set in Appendix D.

In the third step of preprocessing, there are two main types of discrepancies: lost whitespace characters and words sequences, originally separated by line breaks only that were split during sentence split process. The former is solved by inserting the missing whitespace characters, while for the latter we merge the split words into one sequence. The last step of preprocessing is crucial and skipping it will result in the incorrect calculation of the boundary label for augmented text, as the label directly depends on the number of whitespace characters in the text.

After preprocessing, we take a number of consecutive sentences to the left and to the right from the boundary sentence, combine them in a text and determine the label of the boundary word in this new text.

## 4.2 Model Comparison

We used only the transfer learning approach, where a pretrained transformer-based model is fine-tuned on our task in a supervised way. We would like the model to be able to work with long enough sequences because we want the whole human-written prefix to fit in the encoder. Thus, we've determined three models that showed good results on the task of machine-generated text detection:

1. RoBERTa (Liu et al., 2019) has shown good performance in both tasks of boundary detection (Kushnareva et al., 2023) and machine-generated text detection (Macko et al., 2023).
2. Longformer (Beltagy et al., 2020), which was suggested as a baseline by organisers of the task. This model is based on pretrained RoBERTa with novel attention mechanism with a sliding window to long sequences.
3. DeBERTa (He et al., 2021b), which is the state-of-the-art model for machine-generated text detection (Macko et al., 2023). It overcomes the BERT and RoBERTa by introducing a disentangled attention for encoding the position and content of each token separately into two vectors. We decided to test it in the boundary detection task to understand whether it outperform RoBERTa. In our experiments, we fine-tuned DeBERTaV3 (He et al., 2021a) which is an enhanced version of DeBERTa.

All three models are fine-tuned for token classification task. For each token, models predict the

Model	dev	test
RoBERTa-base	9.04 \ 5.78	31.56 \ 30.71
RoBERTa-large	6.72 \ 4.18	25.25 \ 20.66
longformer-base	5.10 \ 5.67	23.16 \ 22.94
longformer-large	4.54 \ 4.40	22.97 \ 20.33
DeBERTaV3-base	3.66 \ 3.15	16.12 \ 13.98
DeBERTaV3-large	2.38 \ 2.54	15.16 \ <b>13.38</b>
Top-1 Submission	-	15.68

Table 1: MAE on original \ augmented dataset and comparison with Top-1 submission on the leaderboard. Longformer-base is suggested as a baseline solution by organisers of competition.

probability of being a boundary token and output the most probable token.

## 5 Experimental Setup

We have used pretrained longformer-4096-base and longformer-4096-large with default hyperparameters to fine-tune Longformer. For experiments with RoBERTa, we have chosen two models: roberta-base and roberta-large. The models were fine-tuned with the set of custom hyperparameters, taken from the original paper (Liu et al., 2019). Finally, for experiments with DeBERTa, we have also chosen two models: deberta-v3-base and deberta-v3-large. The models were fine-tuned with the set of custom hyperparameters taken from He et al. (2021a). All custom hyperparameters are listed in Appendix B.

For all models we set the maximum length of context in tokenizer equal to 512 as there are only few text items in both train and test set with tokenized text length greater than 512. Additionally, we've used the early stopping method for all of our experiments to get rid of epochs dependency.

## 6 Results and Discussion

### 6.1 Main Results

In Tables 1, we compare MAE scores of different models from Section 4.2. There are experiments with models fine-tuned on the original dataset provided by organisers and on the extended dataset with both augmented and original texts.

All models perform better when they are fine-tuned on the extended dataset rather than only on original texts. It clearly shows that even such a simple data augmentation provides a significant boost in performance.

If we compare the results among the models, we will clearly see the dominance of DeBERTaV3 models. For both setups DeBERTaV3-large has shown the best performance and the lowest MAE score on the validation and test datasets. On the setup with the extended dataset, DeBERTaV3-large gets new best MAE score for the competition, which is equal to 13.375. It improves MAE score of the top-1 submission by more than 2 points, from 15.683 to 13.38.

## 6.2 Discussion

Results in Section 6.1 show the importance of both the variety of the data in the dataset and a pretrained model. Leveraging augmented training data significantly increases performance on the task, because it introduces variety in lengths of texts and boundary token positions while preserving the coherence of the texts. We believe that to be the reason why the models perform better when they are fine-tuned on the dataset with augmented data rather than on the original dataset.

Apart from various data, it is also important for the pretrained model to have great generalization capabilities. DeBERTaV3-large has better generalization capabilities than other tested models (He et al., 2021b). The advanced architecture of DeBERTaV3 helps to significantly improve the MAE score in comparison with both RoBERTa and Longformer.

For all three models, large version of each model outperforms base version on both the original and the extended dataset. The reason to this is, greater number of trainable parameters allows models to generalize on training set better.

## 6.3 Error Analysis

We manually inspected texts from the test set on which the best-performing model, DeBERTaV3-large, made serious mistakes. We've limited our inspection set to the texts where the distance between the predicted label and the true label was more than 100 tokens. Thus, we got 276 texts. Only 75 out of these 276 texts were from PeerReview domain, so model made most of its mistakes on the Outfox domain, texts from which are not present in the train set.

These two domains presented in the test set are very different: they vary in the style of formatting, punctuation, and text structuring. The second domain of LLM-generated student essays have a lot of spelling and punctuation problems and it may

confuse the models, as they trained on more formal and literate texts. It would be interesting to evaluate models on each of these domains separately. However, because we do not have domain classification labels in the test set, it is not yet feasible.

## 6.4 Anomalies in Texts

In the majority of texts from the original dataset, the model generates a coherent continuation of the human-written prefix, and it may be hard for a human to guess the boundary word without knowing it. However, there are a number of texts in the data sets that have some flaws in the generated parts.

There are texts in which LLM hallucinated. It either repeated the human prefix or went into a loop where it generated excessive lists with the same beginning. See example of it in Appendix A and in Appendix C. Such hallucinations can be an immediate hint for detector model to put the label boundary near this anomaly. Also, sometimes machine-generated text can have some distinguished features that imply the artificial nature of a particular part of the text. For example, in a number of cases the model begins the generation with the "" (three double quotation marks). It may also be a hint for the detector. A list of other common features we've encountered while examining the test set is provided in Appendix C.

## 7 Conclusion

In this paper we describe the system submitted for SemEval2024-Task 8, the subtask dedicated to hybrid human-machine writing detection. We present a simple yet effective augmentation pipeline. We explore how adding this pipeline to the process of fine-tuning can significantly increase the performance on the task, and provide an analysis of performance of various models with and without our augmentation pipeline. The best model, which is DeBERTaV3-large fine-tuned on a large set of augmented data, receives a new best score according to the leaderboard of the competition. Other fine-tuned models achieve competitive results, ranking in the upper half of the leaderboard and beating the organisers' baseline. As the provided data was limited to English language only, future work might include training multilingual boundary detection solution by mixing training data of different languages and using a multilingual encoder (Chizhikova et al., 2023). Such a system can be used for hybrid AI-writing detection as a standalone solution or can

be integrated into existing NLP frameworks like DeepPavlov (Burtsev et al., 2018).

## References

- Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. [Longformer: The long-document transformer](#).
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS'20*, Red Hook, NY, USA. Curran Associates Inc.
- Mikhail Burtsev, Alexander Seliverstov, Rafael Airapetyan, Mikhail Arkhipov, Dilyara Baymurzina, Nickolay Bushkov, Olga Gureenkova, Taras Khakhulin, Yuri Kuratov, Denis Kuznetsov, and Vasily Konovalov. 2018. [Deeppavlov: An open source library for conversational ai](#). In *NIPS*.
- Anastasia Chizhikova, Vasily Konovalov, and Mikhail Burtsev. 2023. [Multilingual case-insensitive named entity recognition](#). In *Advances in Neural Computation, Machine Learning, and Cognitive Research VI*, pages 448–454, Cham. Springer International Publishing.
- Joseph W. Cutler, Liam Dugan, Shreya Havaldar, and Adam Stein. 2021. [Automatic detection of hybrid human-machine text boundaries](#).
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Liam Dugan, Daphne Ippolito, Arun Kirubarajan, Sherry Shi, and Chris Callison-Burch. 2022. [Real or fake text?: Investigating human ability to detect boundaries between human-written and machine-generated text](#).
- Leon Fröhling and Arkaitz Zubiaga. 2021. [Feature-based detection of automated language models: tackling gpt-2, gpt-3 and grover](#). *PeerJ Computer Science*, 7.
- Sebastian Gehrmann, Hendrik Strobelt, and Alexander Rush. 2019. [GLTR: Statistical detection and visualization of generated text](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 111–116, Florence, Italy. Association for Computational Linguistics.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021a. [Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing](#).
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021b. [Deberta: Decoding-enhanced bert with disentangled attention](#).
- Daphne Ippolito, Daniel Duckworth, Chris Callison-Burch, and Douglas Eck. 2020. [Automatic detection of generated text is easiest when humans are fooled](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1808–1822, Online. Association for Computational Linguistics.
- Ganesh Jawahar, Muhammad Abdul-Mageed, and Laks Lakshmanan, V.S. 2020. [Automatic detection of machine generated text: A critical survey](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2296–2309, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Dongyeop Kang, Waleed Ammar, Bhavana Dalvi, Madeleine van Zuylen, Sebastian Kohlmeier, Eduard Hovy, and Roy Schwartz. 2018. [A dataset of peer reviews \(PeerRead\): Collection, insights and NLP applications](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1647–1661, New Orleans, Louisiana. Association for Computational Linguistics.
- Ryuto Koike, Masahiro Kaneko, and Naoaki Okazaki. 2023. [Outfox: Llm-generated essay detection through in-context learning with adversarially generated examples](#). In *AAAI Conference on Artificial Intelligence*.
- Vasily Konovalov, Oren Melamud, Ron Artstein, and Ido Dagan. 2016. [Collecting Better Training Data using Biased Agent Policies in Negotiation Dialogues](#). In *Proceedings of WOCHAT, the Second Workshop on Chatbots and Conversational Agent Technologies*, Los Angeles. Zerotype.
- Laida Kushnareva, Tatiana Gaintseva, German Magai, Serguei Barannikov, Dmitry Abulkhanov, Kristian Kuznetsov, Irina Piontkovskaya, and Sergey Nikolenko. 2023. [Artificial text boundary detection with topological data analysis and sliding window techniques](#).
- 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](#).

- Yongqiang Ma, Jiawei Liu, Fan Yi, Qikai Cheng, Yong Huang, Wei Lu, and Xiaozhong Liu. 2023. [Ai vs. human – differentiation analysis of scientific content generation](#).
- Dominik Macko, Robert Moro, Adaku Uchendu, Jason Lucas, Michiharu Yamashita, Matúš Pikuliak, Ivan Srba, Thai Le, Dongwon Lee, Jakub Simko, and Maria Bielikova. 2023. [MULTITuDE: Large-scale multilingual machine-generated text detection benchmark](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 9960–9987, Singapore. Association for Computational Linguistics.
- Enrique Manjavacas, Jeroen De Gussem, Walter Daelemans, and Mike Kestemont. 2017. [Assessing the stylistic properties of neurally generated text in authorship attribution](#). In *Proceedings of the Workshop on Stylistic Variation*, pages 116–125, Copenhagen, Denmark. Association for Computational Linguistics.
- Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D. Manning, and Chelsea Finn. 2023. [Detectgpt: zero-shot machine-generated text detection using probability curvature](#). In *Proceedings of the 40th International Conference on Machine Learning*, ICML’23.
- OpenAI. 2023. [Gpt-4 technical report](#).
- Lidiia Ostyakova, Veronika Smilga, Kseniia Petukhova, Maria Molchanova, and Daniel Kornev. 2023. [ChatGPT vs. crowdsourcing vs. experts: Annotating open-domain conversations with speech functions](#). In *Proceedings of the 24th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 242–254, Prague, Czechia. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Conference on Empirical Methods in Natural Language Processing*.
- Irene Solaiman, Miles Brundage, Jack Clark, Amanda Askell, Ariel Herbert-Voss, Jeff Wu, Alec Radford, Gretchen Krueger, Jong Wook Kim, Sarah Kreps, Miles McCain, Alex Newhouse, Jason Blazakis, Kris McGuffie, and Jasmine Wang. 2019. [Release strategies and the social impacts of language models](#).
- Hugo Touvron, Louis Martin, Kevin R. Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Daniel M. Bikel, Lukas Blecher, Cristian Cantón Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony S. Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel M. Kloumann, A. V. Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, R. Subramanian, Xia Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zhengxu Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). *ArXiv*, abs/2307.09288.
- Adaku Uchendu, Thai Le, Kai Shu, and Dongwon Lee. 2020. [Authorship attribution for neural text generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8384–8395, Online. Association for Computational Linguistics.
- Adaku Uchendu, Zeyu Ma, Thai Le, Rui Zhang, and Dongwon Lee. 2021. [TURINGBENCH: A benchmark environment for Turing test in the age of neural text generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2001–2016, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yuxia Wang, Jonibek Mansurov, Petar Ivanov, Jinyan Su, Artem Shelmanov, Akim Tsvigun, Chenxi Whitehouse, Osama Mohammed Afzal, Tarek Mahmoud, Alham Fikri Aji, and Preslav Nakov. 2023. [M4: Multi-generator, multi-domain, and multi-lingual black-box machine-generated text detection](#).
- Yuxia Wang, Jonibek Mansurov, Petar Ivanov, Jinyan Su, Artem Shelmanov, Akim Tsvigun, Chenxi Whitehouse, Osama Mohammed Afzal, Tarek Mahmoud, Giovanni Puccetti, Thomas Arnold, Alham Fikri Aji, Nizar Habash, Iryna Gurevych, and Preslav Nakov. 2024. [Semeval-2024 task 8: Multigenerator, multidomain, and multilingual black-box machine-generated text detection](#). In *Proceedings of the 18th International Workshop on Semantic Evaluation, SemEval 2024*, Mexico, Mexico.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In *Neural Information Processing Systems*.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2019. [Defending against neural fake news](#). Curran Associates Inc., Red Hook, NY, USA.
- Zijie Zeng, Lele Sha, Yuheng Li, Kaixun Yang, Dragan Gašević, and Guanliang Chen. 2023. [Towards automatic boundary detection for human-ai collaborative hybrid essay in education](#).

## A Examples of Texts

Table 2 contains three examples of how authorship change occurs in the texts from the train set. While

the first text contains no signs of a flawed generation, the second and third texts have some flaws. In the second text, the model begins to generate from the capital letter, and its generation is incoherent with the human-written prefix. In the third text, the model starts to repeat the end of the human part, which is a glaring sign of machine-generated text.

Model
I noticed that in Figure 2, the two quantization factors for quantized layers are missing labels. It would be helpful for the reader to understand which layers are being quantized in the figure
Hi Authors, You seem to have submitted two of the same paper? Pls advise Could you please clarify if this is a mistake or if there are any differences between the two submitted papers?
There has been prior work on semi-supervised GAN, though this paper is the first context conditional variant. The novelty of the approach was the novelty of the approach was leveraging in-painting using an adversarial loss to generate contextually relevant images.

Table 2: Examples of texts from train set with different quality of LLM generation and with highlighted human prefix

## B Hyperparameters

For fine-tuning DeBERTaV3 we use hyperparameters, listed by model authors in He et al. (2021a). Table 3 lists these hyperparameters.

Hyperparameters	Large	Base
Optimizer	AdamW	AdamW
Adam $\beta_1, \beta_2$	0.9, 0.999	0.9, 0.999
Adam $\epsilon$	1e-6	1e-6
Warm-up step	50	50
Batch size	4	32
Learning rate (LR)	5e-6	2e-6
Learning rate decay	Linear	Linear
Weight decay	0.01	0.01
Gradient clipping	1.0	1.0

Table 3: Hyperparameters for fine-tuning DeBERTaV3-large and DeBERTaV3-base

Table 4 contains hyperparameters for fine-tuning RoBERTa, also taken from original paper (Liu et al., 2019).

Hyperparameters	values
Optimizer	AdamW
Warm-up steps	50
Batch size	16
Learning rate (LR)	5e-6
Learning rate decay	Linear
Weight decay	0.01

Table 4: Hyperparameters for fine-tuning RoBERTa

## C Examples For Error Analysis

See Table 5.

## D Augmenting Pipeline Scheme

See Figure 2.



Anomaly	Text Id	Example
Excessive repetition	613	...praying for. No more traffic jams, no more parking nightmares, no more car payments, no more insurance, no more maintenance, no more oil changes, no more tire rotations...
Extremely long lists of items	8854	...nice approach to "" + learning skills + learning skills in a sample efficient way + learning skills in an interpretable way + learning skills that can be used on downstream tasks + learning skills that are transferable between domains
JSON-structured hallucinations	5358	...Summary of revisions: * * * """, "title": "Diet Networks: Thin Parameters for Fat Genomics", "abstract": "Learning tasks such as... ...Below are my comments: (1) "" The first sentence of the abstract is too long. It should be divided into two sentences (2) ""
Bizarre formatting	5639	

Table 5: Table with some frequent anomalies in the generated part of texts from test set. The highlighted part is human-written prefix.

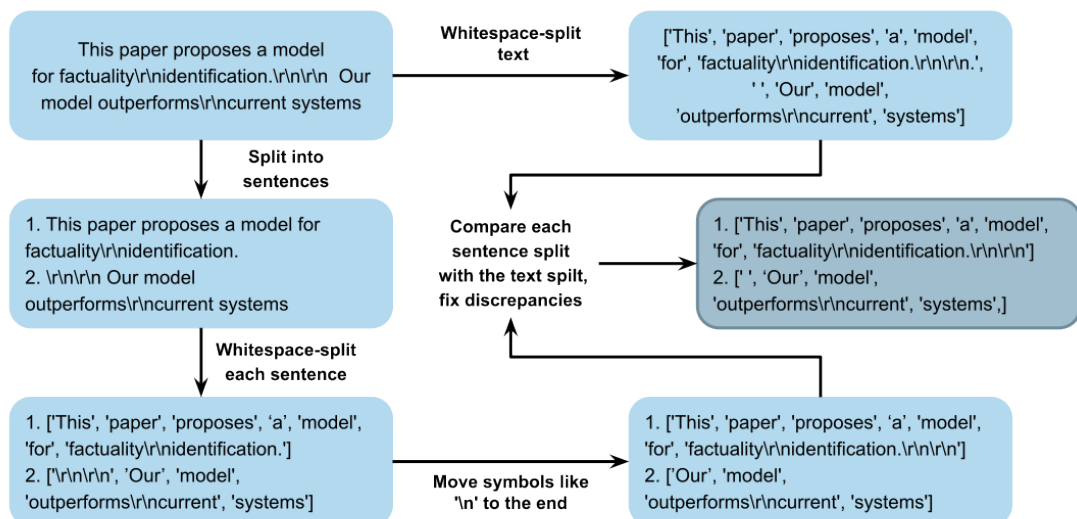


Figure 2: Preprocessing for Augmentation Pipeline