

Advancing LLM detection in the ALTA 2024 Shared Task: Techniques and Analysis

Dima Galat¹,

¹University of Technology Sydney (UTS), Australia,

Correspondence: [dima.galat \[at\] student.uts.edu.au](mailto:dima.galat@student.uts.edu.au)

Abstract

The recent proliferation of AI-generated content has prompted significant interest in developing reliable detection methods. This study explores techniques for identifying AI-generated text through sentence-level evaluation within hybrid articles. Our findings indicate that ChatGPT-3.5 Turbo exhibits distinct, repetitive probability patterns that enable consistent in-domain detection. Empirical tests show that minor textual modifications, such as rewording, have minimal impact on detection accuracy. These results provide valuable insights for advancing AI detection methodologies, offering a pathway toward robust solutions to address the complexities of synthetic text identification.

1 Introduction

The evolution of writing assistants has progressed from simple spell checkers to AI-driven systems (Heidorn, 2000). Advancements of Large Language Models (LLMs), now capable of drafting entire documents, are transforming writing assistants into interactive tools capable of enhancing creativity and productivity (Brown et al., 2020). The introduction of ChatGPT has propelled LLMs into mainstream, quickly gaining them a status of a disruptive technology in many knowledge industries (OpenAI, 2023a). ChatGPT classifier has been discontinued in 2023 seven months after launch citing low accuracy (OpenAI, 2023b).

The analysis of sentiments from early ChatGPT adopters reveals predominantly positive reactions across various domains. However, at the same time concerns have been raised regarding potential misuse and adverse effects in the context of educational activities and news media (Haque et al., 2022). Being able to distinguish between human and machine-generated text is critical for maintaining integrity and transparency in academia, as well as other fields such as journalism.

The rise of human-AI collaborative writing necessitates more advanced detection methods for analysing hybrid texts that incorporate both AI and human-authored sentences. This paper looks at the ALTA 2024 Shared Task challenge (Mollá et al., 2024), where participants develop an automatic detection system to classify sentences in hybrid articles as either human-written or machine-generated. This paper shows ways to improve existing detection methods and promotes more responsible practices in content generation.

2 Background and Related Work

The strategies for a sentence-level detection task predominately focus on the following two approaches: sentence classification, where each sentence in a document is considered independently; or sequence classification, where a document is evaluated as a whole to decide labels for each word and then determine the most frequently-occurring label for the document (Wang et al., 2023). Wang et al. proposed using token-probabilities from different LLMs, aligning local word-wise features to address differences in tokenisation, and then applying convolutions and a linear level for training a sequence classification model exhibiting strong results.

Shi et al. (2024) proposed to look for a boundary between AI and human-authored text, detecting transitions by modeling distances between subsequent sentences in a hybrid document. Experiments demonstrated that this approach consistently improved classification. However, the optimal number of subsequent sentences to be evaluated depends on the document length, and additional considerations are required to account for boundaries that might exist within a hybrid sentence.

Zeng et al. (2024) investigated segmentation within hybrid texts to classify authorship of each segment. The findings suggest to employ a text

segmentation strategy when only a few boundaries exist. Authors note that this is a challenging task; and that short texts provide limited stylistic clues, segment detection is difficult with frequent authorship changes, and human writers are free to select and edit sentences based on their preference (Zeng et al., 2024).

3 Research Methodology

Our goal is to classify sentences generated by ChatGPT-3.5 Turbo mixed with sentences written by humans. We know that at each generation step LLM predicts the next most likely token given the preceding sequence of tokens (i.e. context), or $P(\text{token}_i | \text{context})$. We believe that these marginal probabilities can be used to identify distinct statistical patterns in probability distributions. For example, some high-probability tokens might be favoured by an LLM, whereas human-written text would have higher entropy due to an unexpected choice of words.

3.1 Data and baseline

We were provided a training dataset of 14576 academic and 1500 news articles, containing multiple sentences with a corresponding human/machine label. Validation and test datasets contained 500 and 1000 news articles respectively. When analysing news domain data, we observed that human-written and machine-generated sentences tend to appear in continuous blocks rather than being interwoven or interspersed at the sentence level. Sequences of sentences from each class do not appear to be completely random and follow a pattern resembling the hybrid article generation method of using an LLM with *fill-in task prompts* described by Shi et al..

To rigorously evaluate our sentence classification capabilities without relying on contextual cues, we focus on sentence-level classification. Although a model leveraging entire article context might yield higher accuracy, such approach lies beyond the scope of our current research.

In order to evaluate the importance of domain for building a predictor we have trained baseline models using 3 different versions of the dataset:

1. using all of the training data
2. using random under-sampling of academic articles to match the number of news articles
3. using only news articles

Our baseline model is built using a Naive Bayes classifier, chosen for the efficiency and simplicity. Best results were obtained when using TF-IDF n-gram features up to the length of 5, without stopword filtering (McCallum and Nigam, 1998; Ramos, 2003; Manning et al., 2008). Results suggest that certain phrases and expressions can be favoured by ChatGPT. A relaxed feature independence assumption adds bias and limits the accuracy of the predictions, but this classifier is perfectly suited for comparing statistical properties texts.

We run the model 100 times using different seed selections to account for the random under-sampling of academic articles, and for the variations of a validation data split. The results in Table 1 suggest that that using only news data is sufficient for building a sentence level predictor for this challenge. Moreover, we can see that there are distinctive statistical patterns that can be used to classify these texts.

Dataset	Kappa Score	F1 Weighted
All Data	0.644 ± 0.028	0.83 ± 0.014
Sampling	0.703 ± 0.026	0.86 ± 0.013
Only News	0.716 ± 0.03	0.87 ± 0.014

Table 1: Performance metrics of Naive Bayes for different datasets based on 100 random seed selections, evaluating Cohen’s Kappa Score and F1 Weighted (mean ± standard deviation).

3.2 Model and training

We have attempted using LLM classification zero-shot, however this approach was not getting close to the baseline model. In order to build the best classifier we can, we selected the best base model we could adapt.

LLaMA 3.1 (Meta, 2024), demonstrates a strong capability to generalise across various applications in natural language processing and is very popular in the research community because the release of its weights has facilitated accessibility and further experimentation. Instruction tuning, has emerged as a fine-tuning strategy which augments input-output examples with instructions, enabling instruction-tuned models to generalise more easily to new tasks (Wei et al., 2022).

We are using a model variant with 8 billion parameters which can be trained on a single GPU in a few hours by using a memory-efficient QLORA (Detmers et al., 2023) training approach and 4-bit

quantized weights from Unsloth¹. Our best results are obtained when training in batches of 16 for 3 epochs. This model achieves 0.94 Kappa Score and 0.974 weighted F1 on our validation set, which is significantly above the baseline model results.

4 Results

Overall, it would appear that a 4-bit quantized LLaMA 3.1-8B-Instruct fine-tuned on a domain-specific data can be used to recognise GPT-3.5 Turbo generated content reliably based on the sentence-level evaluation alone. Table 2 shows that our system did well in the competition, outperforming other solutions.

Table 2: Kappa and Accuracy Scores on the test set reported for the participant systems

User	Kappa	Accuracy
our system	0.9320	0.9679
samanjoy2	0.9080	0.9573
lizhuang144	0.8336	0.9235
Qihua	0.7605	0.8914
lewis_math	0.6932	0.8565
dmollaaliod	0.5629	0.7955

5 Discussion

We are still left wondering if our model can reliably detect AI-generated content to prevent misuse. Given the rapid pace of evolution of LLMs, as well as popularisation of generation strategies involving making multiple calls to LLMs, we wonder if our model would still be able to identify ChatGPT-generated sentences if we instructed another model to re-write them.

We are using the same base LLaMA 3.1-8B-Instruct we have fine-tuned for classification to re-write AI-generated sentences in our validation set. Our goal is to see if it will change our sentence classification results. We have used the following prompt, and tried running it up to two times in combination with setting a *temperature* generation parameter to 0.9 to encourage randomness.

Make re-writes to the following sentence without changing the meaning. Only return the sentence, no other information of any kind:\n

Across all of these experiments we obtained good classification results, where the lowest Kappa

¹<https://github.com/unslothai/unsloth>

score produced was 0.89, and weighted F1 was 0.95. This suggests that classification is likely influenced more by the order of certain tokens than by the presence of some specific individual words.

6 Conclusion

Detecting AI-generated content is critical for maintaining authenticity and trust in written communication. We have found that given a small domain-specific corpus a fine-tuned model can reliably identify if a sentence in that corpus has been produced by GPT-3.5 Turbo. Future work could explore how this approach generalises out-of-domain, and to other LLMs.

Different models can end up producing different stylistic features, which means that some day multiple iterations of AI-edits would make it impossible to reliably judge if the text was written by a machine. For now, we have observed that AI-based sentence paraphrasing alone is inadequate to circumvent a classifier trained on in-domain samples. This highlights the importance of efforts involved in developing datasets that accurately represent the behaviour of closed-source models.

References

- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, and et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901.
- Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. 2023. [Qlora: Efficient finetuning of quantized llms](#).
- Mubin Ul Haque, Isuru Dharmadasa, Zarrin Tasnim Sworna, Roshan Namal Rajapakse, and Hussain Ahmad. 2022. [“i think this is the most disruptive technology”](#): Exploring sentiments of chatgpt early adopters using twitter data. *arXiv preprint arXiv:2212.05856*.
- George E. Heidorn. 2000. Intelligent writing assistance. In Ruslan Mitkov, editor, *The Oxford Handbook of Computational Linguistics*, pages 451–465. Oxford University Press.
- Christopher D Manning, Prabhakar Raghavan, and Hinrich Schütze. 2008. *Introduction to Information Retrieval*. Cambridge University Press.
- Andrew McCallum and Kamal Nigam. 1998. A comparison of event models for naive bayes text classification. In *AAAI-98 Workshop on Learning for Text Categorization*, volume 752, pages 41–48.
- Meta. 2024. [The llama 3 herd of models](#).

- Diego Mollá, Qionikai Xu, Zijie Zeng, and Zhuang Li. 2024. Overview of the 2024 alta shared task: Detect automatic ai-generated sentences for human-ai hybrid articles. In *Proceedings of ALTA 2024*.
- OpenAI. 2023a. [Gpt-3.5: Improving language models with instruction tuning](#). Accessed: 2023-02-08.
- OpenAI. 2023b. New ai classifier for indicating ai-written text. <https://openai.com/index/new-ai-classifier-for-indicating-ai-written-text/>. Accessed: 2023-07-20.
- Juan Ramos. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the First International Conference on Machine Learning*, pages 133–142.
- Weijie Shi, Heyang Huang, Yang Xie, Xiang Ren, and Diyi Yang. 2024. Towards detecting ai-generated text within human-ai collaborative hybrid texts. *arXiv preprint arXiv:2403.03506*.
- Pengyu Wang, Linyang Li, Ke Ren, Botian Jiang, Dong Zhang, and Xipeng Qiu. 2023. [Seqxgpt: Sentence-level ai-generated text detection](#). *arXiv preprint arXiv:2310.08903*.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. Finetuned language models are zero-shot learners. In *ICLR 2022 - 10th International Conference on Learning Representations*.
- Zijie Zeng, Shiqi Liu, Lele Sha, Zhuang Li, Kaixun Yang, Sannyuya Liu, Dragan Gašević, and Guanliang Chen. 2024. [Detecting ai-generated sentences in human-ai collaborative hybrid texts: Challenges, strategies, and insights](#). *IJCAI-2024*.