

---

# AuthentiGPT: Detecting Machine-Generated Text via Black-Box Language Models Denoising

---

**Zhen Guo**  
MIT EECS  
Cambridge, MA 02139  
zguo0525@mit.edu

**Shangdi Yu**  
MIT EECS  
Cambridge, MA 02139  
shangdiy@mit.edu

## Abstract

Large language models (LLMs) have opened up enormous opportunities while simultaneously posing ethical dilemmas. One of the major concerns is their ability to create text that closely mimics human writing, which can lead to potential misuse, such as academic misconduct, disinformation, and fraud. To address this problem, we present AuthentiGPT, an efficient classifier that distinguishes between machine-generated and human-written texts. Under the assumption that human-written text resides outside the distribution of machine-generated text, AuthentiGPT leverages a black-box LLM to denoise input text with artificially added noise, and then semantically compares the denoised text with the original to determine if the content is machine-generated. With only one trainable parameter, AuthentiGPT eliminates the need for a large training dataset, watermarking the LLM’s output, or computing the log-likelihood. Importantly, the detection capability of AuthentiGPT can be easily adapted to any generative language model. With a 0.918 AUROC score on a domain-specific dataset, AuthentiGPT demonstrates its effectiveness over other commercial algorithms, highlighting its potential for detecting machine-generated text in academic settings.

## 1 Introduction

Large Language Models (LLMs) have significantly transformed the field of artificial intelligence, offering vast opportunities, but also raising important ethical questions. One of the main issues is their capacity to generate text that closely resembles human writing. Given their extensive scale and accessibility, if misused, LLMs could not only amplify harms such as disinformation and fake news [1], but also undermine academic integrity by assisting unauthorized content generation like essay-writing for students. Such instances highlight the need to differentiate between content generated by machines and that created by humans [2, 3, 4, 5, 6, 7].

In the past, various methods have been developed to tackle the challenge of detecting machine-generated text [8, 9, 10, 11, 12, 13, 14, 15]. One common strategy involves training separate classifiers on datasets containing both human and machine-generated texts with labels [16, 17]. However, the effectiveness of these supervised classifiers often requires a large amount of training data, therefore increasing the cost associated with the training procedure. Another set of evaluation techniques relies on computing the log-likelihood of the language models [18, 19]. However, obtaining the parameters of these models can often be challenging, as they may be inaccessible through chat interfaces or commercial APIs. Watermarking LLM outputs could also offer proactive detection of machine-generated text, maintaining trust and transparency [20, 21, 22]. However, this technique needs a balance between imperceptibility and detectability, while also addressing potential attacks on the watermarks to ensure the integrity of the detection process [23].

Table 1: AUROC scores on the aggregated PubMedQA and GPT generated QA datasets. AuthentiGPT outperforms zero-shot GPT-3.5 and GPT-4, and surpasses GPTZero and Originality.AI.

Methods	AUROC Scores
GPT-3.5 (zero-shot)	0.721
GPT-4 (zero-shot)	0.577
GPTZero	0.797
Originality.AI	0.906
AuthentiGPT	<b>0.918</b>

In this paper, we introduce AuthentiGPT, a novel classification algorithm designed to distinguish between machine-generated and human-written text. Unlike previous methods that rely on extensive labeled data, watermarking, or log-likelihood calculations, AuthentiGPT leverages the language model itself for detection. With the assumption that human-written text resides outside the distribution of machine-generated text [24], the algorithm first introduces synthetic noise to the input text and utilizes a black-box LLM to denoise the text at a noise level. Then, by comparing the denoised text with the original text at the semantic level, a lightweight classifier (with only one free parameter) that is trained on a small set of examples can effectively determine the text was generated by a machine or a human. Our experimental results (Table 1) show that AuthentiGPT outperforms existing methods in detecting machine-generated content within PubMedQA and generated QA datasets, demonstrating the effectiveness of the algorithm. An important advantage of AuthentiGPT is its adaptability. As LLMs continue to improve, AuthentiGPT can be easily adapted to them with minimal effort and modification.

While expanding the capabilities of LLMs, it becomes increasingly crucial to address the ethical considerations surrounding their use. AuthentiGPT represents a valuable advancement in this regard and improves language models’ responsible and ethical application.

## 2 Related Works

### 2.1 Language Model-based Classifiers

Classification tasks can utilize language models by incorporating them with classification layers to perform detections. A conventional classifier that is based on a language model usually operates by passing the language model’s hidden states into a fully-connected layer, which subsequently handles the classification:

$$y = \text{softmax}(W \cdot \text{LM}(S) + b). \tag{1}$$

In this equation,  $\text{LM}(S)$  is the hidden states of the language model for input senetence  $S$ , which has been previously processed by a tokenizer.  $W$  and  $b$  are the parameters of the classifier, and  $y$  is the predicted label. The benefits of using a language model for classification allow the algorithms to capture complex language patterns and context, making them highly effective in understanding and classifying text data. However, depending on the dimension of  $W$ , these classifiers often require a lot of training data and may struggle with the evolving complexity of larger models [25].

### 2.2 Log-likelihood (perplexity) Computation

Perplexity is a measurement in information theory and is commonly used in language modeling to assess how well a probability model predicts a sample [26]. For LLMs, perplexity measures the model’s uncertainty in predicting the next word in a sentence. A lower perplexity indicates that the model is better at predicting the sequence of words, and thus, indicting a higher likelihood of machine-generated content [27]. To compute perplexity, one needs to compute the log-likelihood of text being generated by a model. For input tokens  $X = (x_0, x_1, \dots, x_t)$ , we have

$$\text{PPL}(X) = \exp \left\{ -\frac{1}{t} \sum_i^t \log p_\theta(x_i | x_{<i}) \right\}, \tag{2}$$

where  $\log p_\theta(x_i|x_{<i})$  is the log-likelihood of the  $i$ -th token conditioned on the preceding tokens  $x_{<i}$  given the language model with parameters of  $\theta$ . More sophisticated method, such as DetectGPT, uses the log probability computed by the model with a curvature-based criterion to assess the probability of content being machine-generated [18]. Despite the prospective advantages of this method, it comes with computational demands and requires access to the language model’s parameters, which may not be available in all scenarios. Additionally, its accuracy tends to diminish when faced with complex language models, limiting its overall effectiveness.

### 2.3 Watermark Detection

Watermark detection in machine-generated text involves pinpointing distinctive traits or patterns that suggest its origin or the model that produced it [20]. However, watermark detection encounters obstacles such as the difficulty in recognizing consistent patterns due to the variability of text, safeguarding against adversarial attacks, ensuring watermark robustness without compromising the quality of text, and more. Overall, this field continues to evolve with the advent of more sophisticated watermarks, requiring further development for practical application.

### 2.4 Black-box Detection with N-Gram Analysis

A concurrent study explores the use of a black-box LLM through divergent N-gram analysis (DNA-GPT) for the detection of machine-generated text [28]. They defined the DNA-GPT BScore by:

$$\text{BScore}(S, \Omega) = \frac{1}{K} \sum_{k=1}^K \sum_{n=n_0}^N f(n) \times \frac{|n\text{-grams}(S'_k) \cap n\text{-grams}(S_2)|}{|S'_k| |n\text{-grams}(S_2)|}, \quad (3)$$

where  $S$  is the input text sentence,  $\Omega$  is a set of sentences sampled from the LLM.  $K$  is the sampling repetition,  $f(n)$  is a weight function for different  $n$ -grams,  $S_2$  is a substring of  $S$ , and  $S'_k$  are machine-generated sentences based on  $S \setminus S_2$ . The final classification threshold is determined by balancing the true positive rate and false positive rate. Though the underlining principle of DNA-GPT is similar to AuthentiGPT, AuthentiGPT used an embedding model, instead of n-gram, to extract the features in the input sentences, and utilized a combination of non-linear transformation and unsupervised clustering to determine the threshold, instead of heuristically picking one. The paper claims that the detection strategy is “training-free”, but balancing the true positive rate and false positive rate to determine the classification threshold still requires training examples.

## 3 AuthentiGPT

AuthentiGPT is an efficient detection algorithm that eliminates the requirement for a substantial training dataset, the application of watermarks on the LLM’s output, or the computation of the log-likelihood. It does not need access to the language model’s parameters and can accommodate changes in the LLM with minimal effort. Our algorithm is inspired by [29], which addresses the problem of out-of-distribution detection for images using denoising diffusion probabilistic models (DDPMs). By denoising an input image that has been noised to a range of noise levels, the multi-dimensional reconstruction errors can be obtained and then used to classify out-of-distribution inputs. We apply a similar process to sentences with artificially added noise using an instruction-tuned language model.

On a high-level, AuthentiGPT operates under the assumption that *human-written text resides outside the distribution of machine-generated text*. The algorithm first utilizes a black-box language model to denoise input text with artificially added noise. Then, a semantic comparison is performed between the denoised text and the original text to determine whether it lies within or outside the distribution. In the following, we outline the step-by-step process of our algorithm in Algorithm 1. Inputs to the algorithm are a black-box LLM, some sentences  $S_{\text{test}}$  that we want to classify, some training sentences  $S_{\text{train}}$ , and the labels of these training sentences. It has two parameters,  $\alpha$  (masking ratio) and  $\beta$  (the number of repetitions). We will show that AuthentiGPT can be effective with 10 training samples.

Our algorithm performs the following operations:

---

**Algorithm 1** Detecting Machine-Generated Text

---

```
1: procedure GETSIMILARITY( $S$ , LLM)
2:   for  $i$  in  $[1, 2, \dots, \beta]$  do
3:      $M \leftarrow \text{maskSentences}(S, \alpha)$ 
4:      $D \leftarrow \text{denoiseSentences}(M, \text{LLM})$ 
5:      $E_S \leftarrow \text{computeEmbeddings}(S)$ 
6:      $E_D \leftarrow \text{computeEmbeddings}(D)$ 
7:      $D_{\text{sim},i} \leftarrow \text{getSimilarity}(E_S, E_D)$ 
8:   end for
9:   return  $\text{mean}([D_{\text{sim},1}, \dots, D_{\text{sim},\beta}])$ 
10: end procedure
11: procedure AUTHENTICGPT(LLM,  $S_{\text{test}}$ ,  $S_{\text{train}}$ , labels)
12:    $\mathcal{D}_{\text{train}} \leftarrow \text{GetSimilarity}(S, \text{LLM})$ 
13:    $\text{gm} \leftarrow \text{FindThreshold}(\mathcal{D}_{\text{train}}, \text{labels})$ 
14:    $\mathcal{D}_{\text{sim}} \leftarrow \text{GetSimilarity}(S_{\text{test}}, \text{LLM})$ 
15:   return  $\text{gm.classify}(\mathcal{D}_{\text{sim}})$ 
16: end procedure
```

---

**Algorithm 2** Determine classification threshold

---

```
1: procedure FINDTHRESHOLD( $\mathcal{D}_{\text{sim}}$ , labels)
2:   for  $\lambda$  in  $[\lambda_1, \lambda_2, \dots, \lambda_n]$  do
3:      $\tilde{\mathcal{D}}_\lambda \leftarrow \text{Box-Cox}(\mathcal{D}_{\text{sim}}, \lambda)$ 
4:      $\text{gm}_\lambda = \text{GaussianMixture}(\tilde{\mathcal{D}}_\lambda, \text{n\_class}=2)$ 
5:      $\text{score}_\lambda = \text{AUROC}(\text{gm}_\lambda, \text{labels})$ 
6:   end for
7:   return the  $\text{gm}_\lambda$  and corresponding  $\lambda$  that yield the maximum AUROC score
8: end procedure
```

---

- Randomly masks a portion of the sentences  $S$  determined by a ratio  $\alpha$  to create  $M$ , a version of  $S$  with added noise.
- Denoises sentences  $M$  with the language model with completion or instruction, yielding a denoised version,  $D$ , of the sentences.
- To semantically compare the original and denoised sentences, the algorithm computes embeddings for both  $S$  and  $D$ , denoted as  $E_S$  and  $E_D$ , respectively, and then computes the cosine similarity,  $D_{\text{sim}}$ , between the embeddings of the original and denoised sentences.
- The process repeats  $\beta$  times to allow statistical significance.
- The averaged similarity score is sent to the classifier  $\text{gm}$  for classification. We will explain later in this section how we obtain this classification model.

In our implementation, the mask function is on the word-level with `<unk>` to represent each masked word, the black-box language model LLM is `gpt-3.5-turbo` and embeddings are computed by `text-embedding-ada-002` from OpenAI<sup>1</sup>. Note that the testing sentence  $S_{\text{test}}$  does not necessarily have to be generated by the input LLM; it can be generated by other language models. Averaged cosine similarity serves as a measure of how much the denoised sentences deviate from the original ones after the masking and denoising process. Higher similarity scores indicate that the denoised sentences closely resemble the original ones, thus a higher likelihood of machine-generated text. Conversely, lower similarity scores suggest greater divergence, thus a lower likelihood of machine-generated text. Importantly, computing cosine similarity does not require the parameters of the language model. This is critical because most advanced models are API-based and function as black boxes to users and developers.

To classify the two groups of similarity scores, we use Algorithm 2 to determine the classification boundary. The inputs are  $\mathcal{D}_{\text{sim}}$ , a list of training examples that contain similarities between human-written and machine-generated texts, and their corresponding labels. We use a combination of Box-Cox power transformation [30] and Gaussian Mixture Model (GMM) [31, 32] to perform a

---

<sup>1</sup><https://platform.openai.com/docs>

soft classification. The Box-Cox transformation helps normalize the data, while the GMM enables the identification of the classification boundary in a probabilistic manner. We prefer GMM over other clustering algorithms for its Gaussian distribution assumption, which aligns with our dataset’s characteristics. The output of Algorithm 2 is the Box-Cox parameter  $\lambda$  and its corresponding GMM that yields the maximum AUROC (Area Under the Receiver Operating Characteristic Curve) score [33] on the training set, which will be used to classify the test datasets. Since  $\lambda$  is the only trainable parameter, Algorithm 2 is data-efficient. While our current approach requires minimal training data, transitioning to more advanced classification techniques might be beneficial if more training data is available.




Figure 1: The AUROC scores of AuthentiGPT using different training samples and masking ratios. The x-axis has the number of averaging samples  $\beta$ . Each plot shows the scores using different masking ratios  $\alpha$ .

In our experiments, we evaluated our method ranging from 5 to 20 training examples. Further in the Section 5, we demonstrate that the number of training instances has a minimal impact on the performance of AuthentiGPT. Even with a mere 10 examples from both the original and GPT-generated PubMedQA datasets, the algorithm classifies effectively. This is unsurprising since the only trainable parameter for the algorithm is  $\lambda$  and GMM is an unsupervised clustering model that requires no training. This highlights AuthentiGPT’s remarkable proficiency in differentiating between human-written and machine-generated text. For Algorithm 2, we utilize 100  $\lambda$  values for a grid search. The runtime for this subroutine is negligible when compared to the denoising process.

## 4 Experiment Settings

### 4.1 Dataset

Our dataset includes the original PubMedQA [34] and machine-generated texts from PubMedQA. For evaluation, we use 80 out of 100 instances from each dataset. The rest of the 20 instances from each dataset are combined and used to determine the soft classification threshold.

**Human-written texts** contain 100 original question and answer pairs sourced from PubMedQA. These QA pairs were created by humans and are considered to be reliable and accurate within the biomedical domain.

**Machine-generated texts** include 400 QA pairs that were generated by language models, specifically GPT-3.5 and GPT-4. Two sets of instructions were provided: one to rewrite existing QA pairs and another to generate new QA pairs using the 100 human-written QA pairs as references. These datasets are obtained from [35]. Selected examples are shown in Section A.

### 4.2 Evaluation Metrics

**Accuracy:** It measures the proportion of correct predictions made by a method for a data set that contains only one class (human-generated or machine-generated), considering both true positives and true negatives.

**AUROC:** The ROC curve is created by plotting the true positive rate against the false positive rate at various threshold settings. To compute the AUROC scores for different methods, we aggregate the classification results across all datasets. This approach allows us to obtain a comprehensive assessment of the algorithms’ performance across the entire set of datasets.

### 4.3 Baseline Methods

**GPT-3.5 (Zero-shot):** This method, in a zero-shot setting, generates a binary response ("yes" or "no") indicating whether a given text is machine-generated or not using GPT-3.5.

**GPT-4 (Zero-shot):** Similar to the zero-shot GPT-3.5 method, utilizing GPT-4 model instead.

**GPTZero<sup>2</sup>:** A commercially available classifier, trained to distinguish between human-written and machine-generated texts. It claims to be the most accurate AI detector across use cases.

**Originality.AI<sup>3</sup>:** Another commercial classifier claimed to be the most accurate AI detection tool.

## 5 Results




Figure 2: The figure on the left is the histogram plots for the cosine similarity of the test datasets before the Box-Cox transformation. The figure on the right is the histogram plots for the cosine similarity of the test datasets after the Box-Cox transformation.

GPT-3.5	0.53	0.93	0.80	0.96	0.97
GPT-4	0.19	0.99	0.93	0.95	1.00
GPTZero	0.99	0.51	0.14	0.71	0.74
Originality	0.89	0.85	0.99	0.75	0.97
AuthenticGPT	0.86	0.93	0.53	0.93	0.90
	PubMedQA	GPT3.5-re	GPT4-re	GPT3.5-new	GPT4-new

Figure 3: Accuracy on individual datasets. ‘GPT3.5-re’ and ‘GPT4-re’ are the datasets obtained by re-writing using GPT-3.5 and GPT-4 respectively. ‘GPT3.5-new’ and ‘GPT4-new’ are the datasets obtained by generating new QA pairs using GPT-3.5 and GPT-4 respectively.

In Figure 1, we present the AUROC scores of AuthenticGPT on the combined PubMedQA and GPT-generated QA dataset with different masking ratios and the number of training samples. As the masking ratio  $\alpha$  increases, the performance of the algorithm also improves, reaching a plateau at around a masking ratio of 0.08. Intuitively, finding an optimal masking ratio requires meticulous tuning. A high masking ratio may result in a loss of crucial semantic information and thus limits the context for the LLM to work with, making the denoising task challenging. Consequently, the denoised sentences, whether machine-generated or human-written, would deviate significantly in semantics

<sup>2</sup><https://gptzero.me/>

<sup>3</sup><https://originality.ai/>

from their original versions regardless, making the classification task difficult. On the other hand, if the masking ratio is very small, the denoising task would be trivial. As a result, denoised sentences may remain semantically close to the original ones, which are also difficult to classify. Therefore, these considerations necessitate careful balance in the masking ratio for optimal performance for the algorithm.

On the other hand, the influence of the number of training samples appears minimal except for  $\alpha = 0.04$ . Again, this is unsurprising since the only trainable parameter for the algorithm is  $\lambda$  and GMM is an unsupervised clustering model that requires no training. A higher  $\beta$  consistently improves the performance. In other words, more samples increase the statistical significance of the classification results. However, it's crucial to note that the runtime for the algorithm is linearly correlated with  $\beta$ . In the rest of the experiments, we will use a masking ratio ( $\alpha$ ) of 0.08, number of averaging sample ( $\beta$ ) of 10, and 20 training samples for AuthentiGPT.

In Figure 2, we show the histogram of the cosine similarity between the original text and the denoised text from the black-box language model `gpt-3.5-turbo` with  $\alpha$  of 0.08 and  $\beta$  of 10. The hyperparameter  $\lambda$  is of 0.5, which will also be used in Figure 3. After applying the Box-Cox transformation, the dataset distribution resembles a Gaussian distribution, and the separation between different dataset distributions becomes more discernible. This statistical technique is crucial for our final classification using GMM clustering.

Table 1 shows the AUROC scores for various detection methods on the combined PubMedQA and GPT-generated QA dataset. Notably, AuthentiGPT demonstrates superior performance compared to zero-shot GPT-3.5 and GPT-4, and slightly surpasses Originality.AI, a high-end commercial classifier. GPTZero, another commercial classifier, exhibits intermediate performance.

In Figure 3, we present a comparative analysis of the methods on each individual dataset. The performance is evaluated using the accuracy metric. We observe that some methods exhibit a tendency to classify sentences into one class. Specifically, GPT-3.5 and GPT-4 tend to classify sentences as machine-generated while GPTZero tends to classify sentences as human-written. As a result, though GPT-3.5 and GPT-4 correctly classify the machine-generated sentences, they fail at correctly classifying human-written sentences (original PubMedQA). Similarly, though GPTZero correctly classifies human-written sentences, it has much lower accuracy in classifying machine-generated sentences. Across tasks, Originality.AI is consistently good except for `gpt3.5-new` task. AuthentiGPT performs well across the board, with the exception in `gpt4_rewrite`. These variations can be attributed to the inherent complexities and subtle distinctions within each dataset. For GPTZero and AuthentiGPT, classifying GPT rewrites is more challenging than new QA tasks, potentially due to the retention of the human-written text style and semantics in the rewrite datasets [12]. While tasks involving generated new QAs are comparatively easier as the machine-generated patterns become more distinguishable.

## 6 Limitations and Future Work

Although AuthentiGPT outperforms all other methods on the aggregated PubMedQA datasets, it exhibits shortcomings in individual datasets, particularly in `gpt4_rewrite`. In addition, several limitations must be acknowledged. Firstly, the extent to which the algorithm's effectiveness applies beyond the biomedical field is unclear. Future work should focus on assessing its generalizability across a variety of domains. Secondly, our assumption that human-written text lies outside the distribution of machine-generated text may not always hold true. This could potentially impact the algorithm's performance. Finally, the risk of false positives in detection, which could undermine trust in assessment tools, is a critical but unaddressed concern. These considerations highlight the need for further optimization of the method, such as including additional black-box language models for ensemble averaging or using neural classifiers with more training samples.

## 7 Conclusion

In conclusion, AuthentiGPT offers a novel approach for detecting machine-generated text using small amount of training data. As language models continue to evolve and advance, the importance of tools like AuthentiGPT would become increasingly apparent. Those tools may play a crucial role in ensuring the ethical and responsible application of these models.

## References

- [1] Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. Defending against neural fake news. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [2] Sebastian Gehrmann, Hendrik Strobelt, and Alexander M. Rush. Gltr: Statistical detection and visualization of generated text, 2019.
- [3] Daphne Ippolito, Daniel Duckworth, Chris Callison-Burch, and Douglas Eck. 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, July 2020. Association for Computational Linguistics.
- [4] Evan Crothers, Nathalie Japkowicz, and Herna Viktor. Machine generated text: A comprehensive survey of threat models and detection methods. *arXiv preprint arXiv:2210.07321*, 2022.
- [5] Yogesh K Dwivedi, Nir Kshetri, Laurie Hughes, Emma Louise Slade, Anand Jeyaraj, Arpan Kumar Kar, Abdullah M Baabdullah, Alex Koohang, Vishnupriya Raghavan, Manju Ahuja, et al. “so what if chatgpt wrote it?” multidisciplinary perspectives on opportunities, challenges and implications of generative conversational ai for research, practice and policy. *International Journal of Information Management*, 71:102642, 2023.
- [6] Michael Liebrez, Roman Schleifer, Anna Buadze, Dinesh Bhugra, and Alexander Smith. Generating scholarly content with chatgpt: ethical challenges for medical publishing. *The Lancet Digital Health*, 5(3):e105–e106, 2023.
- [7] Xinlei He, Xinyue Shen, Zeyuan Chen, Michael Backes, and Yang Zhang. Mgtbench: Benchmarking machine-generated text detection. *arXiv preprint arXiv:2303.14822*, 2023.
- [8] Sameer Badaskar, Sachin Agarwal, and Shilpa Arora. Identifying real or fake articles: Towards better language modeling. In *Proceedings of the Third International Joint Conference on Natural Language Processing: Volume-II*, 2008.
- [9] Daria Beresneva. Computer-generated text detection using machine learning: A systematic review. In *Natural Language Processing and Information Systems: 21st International Conference on Applications of Natural Language to Information Systems, NLDB 2016, Salford, UK, June 22-24, 2016, Proceedings 21*, pages 421–426. Springer, 2016.
- [10] Kanish Shah, Henil Patel, Devanshi Sanghvi, and Manan Shah. A comparative analysis of logistic regression, random forest and knn models for the text classification. *Augmented Human Research*, 5:1–16, 2020.
- [11] Ganesh Jawahar, Muhammad Abdul-Mageed, and Laks VS Lakshmanan. Automatic detection of machine generated text: A critical survey. *arXiv preprint arXiv:2011.01314*, 2020.
- [12] Sandra Mitrović, Davide Andreoletti, and Omran Ayoub. Chatgpt or human? detect and explain. explaining decisions of machine learning model for detecting short chatgpt-generated text. *arXiv preprint arXiv:2301.13852*, 2023.
- [13] Vinu Sankar Sadasivan, Aounon Kumar, Sriram Balasubramanian, Wenxiao Wang, and Soheil Feizi. Can ai-generated text be reliably detected? *arXiv preprint arXiv:2303.11156*, 2023.
- [14] Souradip Chakraborty, Amrit Singh Bedi, Sicheng Zhu, Bang An, Dinesh Manocha, and Furong Huang. On the possibilities of ai-generated text detection. *arXiv preprint arXiv:2304.04736*, 2023.
- [15] Ruixiang Tang, Yu-Neng Chuang, and Xia Hu. The science of detecting llm-generated texts. *arXiv preprint arXiv:2303.07205*, 2023.
- [16] Fabrizio Sebastiani. Machine learning in automated text categorization. *ACM computing surveys (CSUR)*, 34(1):1–47, 2002.



- [17] Niful Islam, Debopom Sutradhar, Humaira Noor, Jarin Tasnim Raya, Monowara Tabassum Maisha, and Dewan Md Farid. Distinguishing human generated text from chatgpt generated text using machine learning. *arXiv preprint arXiv:2306.01761*, 2023.
- [18] Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D Manning, and Chelsea Finn. Detectgpt: Zero-shot machine-generated text detection using probability curvature. *arXiv preprint arXiv:2301.11305*, 2023.
- [19] Zhijie Deng, Hongcheng Gao, Yibo Miao, and Hao Zhang. Efficient detection of llm-generated texts with a bayesian surrogate model. *arXiv preprint arXiv:2305.16617*, 2023.
- [20] John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. A watermark for large language models. *arXiv preprint arXiv:2301.10226*, 2023.
- [21] John Kirchenbauer, Jonas Geiping, Yuxin Wen, Manli Shu, Khalid Saifullah, Kezhi Kong, Kasun Fernando, Aniruddha Saha, Micah Goldblum, and Tom Goldstein. On the reliability of watermarks for large language models, 2023.
- [22] Travis Munyer and Xin Zhong. Deeptextmark: Deep learning based text watermarking for detection of large language model generated text. *arXiv preprint arXiv:2305.05773*, 2023.
- [23] Kalpesh Krishna, Yixiao Song, Marzena Karpinska, John Wieting, and Mohit Iyyer. Paraphrasing evades detectors of ai-generated text, but retrieval is an effective defense. *arXiv preprint arXiv:2303.13408*, 2023.
- [24] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations*, 2020.
- [25] Weixin Liang, Mert Yuksekgonul, Yining Mao, Eric Wu, and James Zou. Gpt detectors are biased against non-native english writers. *arXiv preprint arXiv:2304.02819*, 2023.
- [26] Robert C Moore and William Lewis. Intelligent selection of language model training data. In *Proceedings of the ACL 2010 conference short papers*, pages 220–224, 2010.
- [27] Christoforos Vasilatos, Manaar Alam, Talal Rahwan, Yasir Zaki, and Michail Maniatakos. Howkgpt: Investigating the detection of chatgpt-generated university student homework through context-aware perplexity analysis. *arXiv preprint arXiv:2305.18226*, 2023.
- [28] Xianjun Yang, Wei Cheng, Linda Petzold, William Yang Wang, and Haifeng Chen. Dna-gpt: Divergent n-gram analysis for training-free detection of gpt-generated text. *arXiv preprint arXiv:2305.17359*, 2023.
- [29] Mark S. Graham, Walter H.L. Pinaya, Petru-Daniel Tudosiu, Parashkev Nachev, Sebastien Ourselin, and Jorge Cardoso. Denoising diffusion models for out-of-distribution detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2947–2956, June 2023.
- [30] Remi M Sakia. The box-cox transformation technique: a review. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 41(2):169–178, 1992.
- [31] Carl Rasmussen. The infinite gaussian mixture model. *Advances in neural information processing systems*, 12, 1999.
- [32] Douglas A Reynolds et al. Gaussian mixture models. *Encyclopedia of biometrics*, 741(659-663), 2009.
- [33] Elizabeth R DeLong, David M DeLong, and Daniel L Clarke-Pearson. Comparing the areas under two or more correlated receiver operating characteristic curves: a nonparametric approach. *Biometrics*, pages 837–845, 1988.
- [34] Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William W Cohen, and Xinghua Lu. Pubmedqa: A dataset for biomedical research question answering. *arXiv preprint arXiv:1909.06146*, 2019.
- [35] Zhen Guo, Peiqi Wang, Yanwei Wang, and Shangdi Yu. Dr. llama: Improving small language models in domain-specific qa via generative data augmentation. *arXiv preprint arXiv:2305.07804*, 2023.

## A Examples from our datasets

In this section, we show one example from each of our five datasets. The PubMedQA is included in prompting language models to rewrite or generate a new question-answering pair in the biomedical domain.

PubMedQA

**Question:** Syncope during bathing in infants, a pediatric form of water-induced urticaria? context: Apparent life-threatening events in infants are a difficult and frequent problem in pediatric practice. The prognosis is uncertain because of risk of sudden infant death syndrome. Eight infants aged 2 to 15 months were admitted during a period of 6 years; they suffered from similar maladies in the bath: on immersion, they became pale, hypotonic, still and unreactive; recovery took a few seconds after withdrawal from the bath and stimulation. Two diagnoses were initially considered: seizure or gastroesophageal reflux but this was doubtful. The hypothesis of an equivalent of aquagenic urticaria was then considered; as for patients with this disease, each infant's family contained members suffering from dermatographism, maladies or eruption after exposure to water or sun. All six infants had dermatographism. We found an increase in blood histamine levels after a trial bath in the two infants tested. The evolution of these "aquagenic maladies" was favourable after a few weeks without baths. After a 2-7 year follow-up, three out of seven infants continue to suffer from troubles associated with sun or water. answer: "Aquagenic maladies" could be a pediatric form of the aquagenic urticaria.

**Answer:** the answer to the question given the context is yes

gpt3.5-rewrite

**Question:** Can the use of a new medication reduce symptoms of depression in adults with major depressive disorder? Context: Major depressive disorder (MDD) is a common mental health condition that can significantly affect a person's quality of life. Current treatments for MDD include therapy and medication, but many patients do not respond well to these options or experience side effects. A new medication, called WXY-19, has been developed that targets a specific neurochemical pathway involved in depression. A clinical trial was conducted with 100 adults diagnosed with MDD. Participants were randomly assigned to receive either WXY-19 or a placebo for 12 weeks. Results showed that participants who received WXY-19 had a significant reduction in depressive symptoms compared to those who received the placebo.

**Answer:** Yes, the use of WXY-19 can reduce symptoms of depression in adults with major depressive disorder.

gpt4-rewrite

**Question:** Is syncope during bathing in infants a potential pediatric form of water-induced urticaria? context: Apparent life-threatening events in infants pose a challenging and common issue in pediatric practice, with an uncertain prognosis due to sudden infant death syndrome risks. Over six years, eight infants aged 2 to 15 months experienced similar symptoms while bathing: they became pale, hypotonic, motionless, and unresponsive upon immersion, recovering within seconds after being removed from the bath and stimulated. Initial diagnoses considered were seizures or gastroesophageal reflux, but these were doubtful. The possibility of a pediatric equivalent to aquagenic urticaria was then explored; like patients with this condition, each infant's family had members with dermatographism, ailments, or rashes triggered by water or sun exposure. All six infants demonstrated dermatographism. In two tested infants, blood histamine levels increased after a trial bath. The progression of these "aquagenic ailments" improved after a few weeks without bathing. Following a 2-7 year follow-up, three of the seven infants continued to experience issues related to sun or water exposure.

**Answer:** Based on the context, it appears that "aquagenic ailments" might be a pediatric form of water-induced urticaria, so the answer to the question is yes.

gpt3.5-new

**New question:** Is there a correlation between maternal smoking during pregnancy and the risk of sudden infant death syndrome (SIDS)?  
**Context:** Sudden infant death syndrome (SIDS) is the unexpected and sudden death of an apparently healthy infant, typically during sleep. It is the leading cause of death in infants between 1 month and 1 year of age. Several factors have been identified as risk factors for SIDS, including maternal smoking during pregnancy. Maternal smoking is known to have negative effects on fetal development, and it has been suggested that it may also increase the risk of SIDS.

**Answer:** Yes, maternal smoking during pregnancy has been identified as a risk factor for sudden infant death syndrome (SIDS).

gpt4-new

**New question:** Can genetic predisposition to dermographism play a role in infants experiencing syncope during bathing?  
**context:** A study investigated eight infants aged 2 to 15 months who were admitted over a period of 6 years with similar symptoms during bathing. These infants became pale, hypotonic, still, and unreactive upon immersion in water, and recovered within seconds once removed from the bath and stimulated. Initial diagnoses considered were seizure or gastroesophageal reflux, but these were deemed unlikely. Researchers then considered the possibility of a pediatric form of aquagenic urticaria. Each infant's family had members with dermographism, skin reactions to water or sun exposure, and all six infants tested had dermographism. An increase in blood histamine levels after a trial bath was observed in two infants tested. After a few weeks without baths, the condition improved, but after a follow-up of 2-7 years, three out of seven infants continued to have issues related to sun or water exposure.

**Answer:** Yes, genetic predisposition to dermographism may play a role in infants experiencing syncope during bathing, as it may be related to a pediatric form of aquagenic urticaria.

## Ethical Statement

Our research on AuthentiGPT strictly adhered to ethical standards, with the objective of promoting safety and trust in the AI community. We recognize that AuthentiGPT does not solve all ethical challenges related to LLMs, advocating for ongoing interdisciplinary collaboration. Our commitment is to maintain ethical transparency and to constantly reassess our ethical considerations as the field of AI continues to evolve.

## Acknowledgement

We would like to thank OpenAI, GPTZero, and Originality.ai for providing their API services. We are grateful to Francis Huang for providing helpful feedback and discussion.