

Detecting ChatGPT: A Survey of the State of Detecting ChatGPT-Generated Text

Mahdi Dhaini

Wessel Poelman

Ege Erdogan

Technical University of Munich
Department of Computer Science
Germany

{mahdi.dhaini, wessel.poelman, ege.erdogan}@tum.de

Abstract

While recent advancements in the capabilities and widespread accessibility of generative language models, such as ChatGPT (OpenAI, 2022), have brought about various benefits by generating fluent human-like text, the task of distinguishing between human- and large language model (LLM) generated text has emerged as a crucial problem. These models can potentially deceive by generating artificial text that appears to be human-generated. This issue is particularly significant in domains such as law, education, and science, where ensuring the integrity of text is of the utmost importance. This survey provides an overview of the current approaches employed to differentiate between texts generated by humans and ChatGPT. We present an account of the different datasets constructed for detecting ChatGPT-generated text, the various methods utilized, what qualitative analyses into the characteristics of human versus ChatGPT-generated text have been performed, and finally, summarize our findings into general insights.

1 Introduction

LLMs have been showing remarkable abilities in generating fluent, grammatical, and convincing text. The introduction of ChatGPT (OpenAI, 2022) has been widely regarded as a significant and controversial milestone for LLMs. Models such as GPT-3 (Brown et al., 2020) and PaLM (Chowdhery et al., 2022) already demonstrated the power of LLMs in many natural language processing (NLP) tasks. ChatGPT is the first model that has seen widespread adoption outside NLP research.

The increased performance of LLMs raises important questions regarding their potential societal impact. The risks of LLMs are numerous, from confidently presenting false information to generating fake news on a large scale (Sheng et al., 2021;

Weidinger et al., 2022). ChatGPT is no exception in this regard (Zhuo et al., 2023).

Instances of the misuse of ChatGPT have already been documented in various domains, including education (Cotton et al., 2023), scientific writing (Gao et al., 2022), and the medical field (Anderson et al., 2023). Given this context, the detection of machine-generated text is gaining considerable attention. This detection is part of a larger push towards responsible and appropriate usage of generative language models (Kumar et al., 2023).

In addition to academic interest, a growing number of commercial parties are trying to solve this task. Recent work from Pegoraro et al. (2023) gives an overview of commercial and freely available online tools. They come close to the current work. However, we limit our scope to academic work and provide additional background information on methods, datasets, and qualitative insights.

Many approaches, datasets and shared tasks¹ have been put forth recently to tackle the *general* (i.e., not specific for ChatGPT) task of detecting machine-generated text (Jawahar et al., 2020). Given the enormous use and cultural impact of ChatGPT, we limit our review to datasets and methods developed directly for ChatGPT. We discuss these methods in the context of the controversial position ChatGPT is in, namely that it is a closed-source system with very little information available regarding its training setup or model architecture at the time of writing. We outline what general methods exist for this task and review recent work that directly focuses on datasets and methods for ChatGPT.

Given the peculiar place ChatGPT is in, we also consolidate qualitative insights and findings from the works we discuss that might help humans to detect ChatGPT-generated text. These include lin-

¹For instance, AuTexTification or CLIN33 Shared Task.

guistic features or writing styles to look out for. Lastly, we present outstanding challenges for this detection task and possible future directions.

2 Related Work on Detecting Machine-Generated Text

LLMs have become a driving force in many language processing-related benchmarks and tasks (Radford et al., 2018; Brown et al., 2020; Chowdhery et al., 2022). LLMs can solve complex NLP tasks and generate convincing and trustworthy-looking text. However, they are also prone to generating false and misleading information, generally referred to as *hallucinating* (Lin et al., 2022). Additionally, misuse of these models can pose significant risks in academia, journalism, and many other areas. Currently, human judges are decent at spotting machine-generated text from older LLMs such as GPT-2 (Ippolito et al., 2020; Dugan et al., 2020, 2023). Still, the increasing abilities of LLMs give rise to the need for more sophisticated detection tools and models.

A recent survey by Crothers et al. (2023) provides a thorough overview of risks, approaches, and detection methods. They discuss interesting aspects such as the effect of domains on the detection task, adversarial attacks, and societal impacts of generated texts. Work done by Jawahar et al. (2020) inspects the field of machine-generated text detection. It outlines three main detection methods: a classifier trained from scratch, zero-shot detection by a language model, and a fine-tuned language model as a classifier. Recently, detection methods from computer vision have also been tried on language models, such as watermarking (Kirchenbauer et al., 2023a,b) or trying to find model-identifying artifacts in generated content (Tay et al., 2020). To use and evaluate these methods, fine-grained access to the source model is required in training and inference time. Both these preconditions are not the case with ChatGPT at the time of writing.

When discussing detection methods, an essential factor to consider is access to the log probability output of a model. This is the probability distribution over the vocabulary of a model for the next token to be generated. Numerous successful detection methods evaluate the average log probability per token combined with a threshold in a zero-shot setting (Gehrmann et al., 2019; Ippolito et al., 2020; Mitchell et al., 2023). This method is model agnostic and generally performs quite well. At the

time of writing, users of ChatGPT do not have access to these probabilities. Without this access or knowledge about model internals, detection methods are limited to using just the generated text in a binary classification setting, with the options being *human* or *machine*. These methods use simple classifiers trained on n-grams (Solaiman et al., 2019; Ippolito et al., 2020) or fine-tuned pre-trained language models (Uchendu et al., 2020; Ippolito et al., 2020; Zellers et al., 2020).

Another group of detection tools we want to mention are the *human-machine collaboration systems*, as Jawahar et al. (2020) labels them. These tools do not necessarily classify a passage directly but assist a human in making that decision. The previously mentioned work by Gehrmann et al. (2019) visualizes tokens in different colors, depending on where a given token ends up in the top- k most probable tokens from the model. This can also assist a human judge in spotting which part of a larger text might be machine-generated, such as possibly rephrased or copied sections for example. As mentioned, this method requires access to output probabilities, so it is not usable for ChatGPT. Another tool to help humans in the detection task is to outline the linguistic properties and characteristics of machine-generated text. This was one of the main goals of the *Real or Fake Text?* (RoFT) game created by Dugan et al. (2020, 2023). This game asked players to decide if a machine partially wrote a piece of text, and if yes, where the transition point from human to machine is in the text. This resulted in a considerable dataset of annotations and indicators humans look for in detecting machine-generated text.

Another area of research that might help humans to make this decision is explainable AI. As we will see, some papers we discuss use explainability methods, such as SHAP (Lundberg and Lee, 2017), in their approaches. These methods help to better understand how detectors make their predictions. Such methods can help provide insights on the input features that most contribute to a prediction, which, in turn, can facilitate analyses of the differences between human and ChatGPT writing styles.

As far as we know, the previously mentioned work by Crothers et al. (2023) and Jawahar et al. (2020) come closest to ours. They discuss detection methods and datasets but not ChatGPT. The work from Pegoraro et al. (2023) does mention ChatGPT, among other models, but focuses mainly on online

detection tools.

Our contributions are the following:

- We provide an overview of *general* approaches to machine-generated text detection.
- We outline research that specifically addresses the detection of ChatGPT-generated text and how this relates to the general approaches.
- We show the datasets that are created and used for this detection task.
- We summarize the qualitative analyses that these recent works provide and try to give general insights.

3 Review of Approaches for Detecting ChatGPT-Generated Text

3.1 Datasets

Table 1 shows datasets that can be used to perform analyses or train models to distinguish between human and ChatGPT written text. We describe how they were collected and provide further information on their domains and public availability.

3.1.1 Guo et al. 2023 (HC3)

Available in both Chinese and English, the Human ChatGPT Comparison Corpus (HC3) contains question-answer pairs collected from different datasets such as OpenQA (Yang et al., 2015) and Reddit ELI5 (Fan et al., 2019). These questions are then given to ChatGPT with context-sensitive prompts (e.g., asking ChatGPT to answer *like I am five* for the Reddit ELI5 dataset) so that each question has one human-generated and one ChatGPT-generated answer.

3.1.2 Yu et al. 2023 (CHEAT)

The ChatGPT-written Abstract (CHEAT) dataset contains human- and ChatGPT-generated title-abstract pairs for computer science papers, with the titles and human-written abstracts fetched from IEEE Xplore. Artificial abstracts are generated in three ways:

- *Generate*: ChatGPT is directly prompted to write an abstract given the title and keywords.
- *Polish*: ChatGPT is given human-written abstracts and is told to “polish” them.
- *Mix*: Text from human-written and polished abstracts are mixed at the sentence level.

The CHEAT dataset also covers adversarial scenarios as the *Polish* and *Mix* methods correspond to methods a malicious user might try to evade detection.

3.1.3 He et al. 2023 (MGTBench)

The Machine Generated Text Benchmark (MGT-Bench) uses three question-answering datasets: TruthfulQA (Lin et al., 2022), SQuAD1 (Rajpurkar et al., 2016), and NarrativeQA (Kočíský et al., 2018). Questions are randomly sampled from each dataset, and ChatGPT is prompted to answer them with the appropriate context (e.g., with a relevant passage and instructions for NarrativeQA).

Although our primary focus is ChatGPT, MGT-Bench contains text generated by different language models and thus can be used to benchmark detection methods across models.

3.1.4 Liu et al. 2023 (ArguGPT)

The ArguGPT dataset contains prompts and responses from various English learning corpora, such as WECCL (Zhi-jia, 2008), TOEFL11 (Blanchard et al., 2013), and hand-picked from graduate record examinations (GRE) preparation material. The texts are from essay writing assignments about a given topic or standpoint. GPT models are prompted to write responses, but their output is processed for grammatical errors and to remove obvious signs of ChatGPT-generated text (e.g., “As a large language model. . .”).

3.1.5 Vasilatos et al. 2023

The dataset used in Vasilatos et al. (2023) for detection builds on Ibrahim et al. (2023), a dataset of questions with metadata and student answers from various university courses. ChatGPT is directly prompted with the questions three times to obtain three human and ChatGPT answers for each question.

3.1.6 Mitrović et al. 2023

Attempting to build a classifier to detect ChatGPT-generated restaurant reviews, Mitrović et al. (2023) build on the Kaggle restaurant reviews dataset² and prompt ChatGPT to generate reviews of various kinds (e.g., “write a review for a bad restaurant”). Additionally, ChatGPT is prompted to rephrase the human-written reviews to create an adversarial set.

²<https://www.kaggle.com/competitions/restaurant-reviews/overview>

Dataset (name)	Domain	Public	OOD	Size and Setup
Guo et al. 2023 (HC3-English)	Multi-domain	✓	×	Q&A Questions: 24,322 Human-A: 58,546 ChatGPT-A: 26,903
Guo et al. 2023 (HC3-Chinese)	Multi-domain	✓	×	Q&A Questions: 12,853 Human-A: 22,259 ChatGPT-A: 17,522
Yu et al. 2023 (CHEAT)	Scientific	×	✓	Abstracts Human: 15,395 ChatGPT: 35,304
He et al. 2023 (MGTBench)	General	✓	×	Q&A pairs Human: 2,817 ChatGPT: 2,817
Liu et al. 2023 (ArguGPT)	Education	✓	×	Essays Human: 4,115 ChatGPT: 4,038
Vasilatos et al. 2023	Education	Human*	×	Q&A Questions: 320 Human-A: 960 ChatGPT-A: 960
Mitrović et al. 2023	General	Human*	✓	Reviews Human: 1,000 ChatGPT-query: 395 ChatGPT-rephrase: 1,000
Weng et al. 2023	Scientific	Human	×	Title-Abstract pairs Human: 59,232 ChatGPT: 59,232
Antoun et al. 2023a	General	✓	✓	Q&A HC3-English OOD-ChatGPT: 5,969
Liao et al. 2023	Medical	Human	×	Abstracts and records Human: 2,200 ChatGPT: 2,200

Table 1: Datasets used in ChatGPT-generated text detection, with public availability information (if a dataset is available, it can be accessed by clicking on its *Public* column entry). The *Human* entry in the *Public* column signals that only human-written text datasets are made public. The *OOD* (out-of-domain) column signals if a dataset contains examples generated in a different way than the main part (e.g., rephrasing of human-written text). *Authors state it will be made available at a future date.

3.1.7 Weng et al. 2023

Weng et al. (2023) expand on Narechania et al. (2022)’s dataset of title-abstract pairs fetched from top data visualization venues by prompting ChatGPT to write abstracts given the titles. Compared

to another dataset of title-abstract pairs, CHEAT (Yu et al., 2023), Weng et al. (2023)’s dataset contains more examples but lacks the adversarial samples included in CHEAT.

3.1.8 Antoun et al. 2023a

Antoun et al. (2023a) extend HC3 (Guo et al., 2023) by translating its English part to French using Google Translate and add further French out-of-domain (OOD) examples to make models trained on this data more robust. The OOD dataset consists of direct French responses by ChatGPT and BingChat to translated questions from the HC3 dataset (as opposed to translating the answers as done originally), question-answer pairs from the French part of the multi-lingual QA dataset MFAQ (De Bruyn et al., 2021), and sentences from the French Treebank dataset (Le Monde corpus). Finally, the dataset also contains a small number of adversarial examples written by humans with access to ChatGPT to obtain a similar style to that of ChatGPT.

3.1.9 Liao et al. 2023

Focusing on the medical domain, Liao et al. (2023) build on two public medical datasets: a set of medical abstracts from Kaggle³ and radiology reports from the MIMIC-III dataset (Johnson et al., 2016). ChatGPT is given parts of an example medical abstract or a radiology report for the machine-generated samples and is prompted to continue writing it. The authors state that text continuation can generate more human-like text compared to rephrasing or direct prompting.

3.2 Methods

In this section, we report on the various methods proposed for detecting ChatGPT-generated text. The scope of this review does not include the evaluation or comparison of the results obtained from these methods. This limitation primarily arises from the absence of a common experimental setup and the utilization of different datasets and metrics. Table 2 provides an overview of these recent approaches.

Some previous works have utilized transformer-based models to classify text generated by ChatGPT and human-written text, as demonstrated by Mitrović et al. (2023). Their approach consists of two components: a detection model and a framework to explain the decisions made by this model. They first fine-tune an uncased version of DistilBERT (Sanh et al., 2019) and then employ SHAP to provide local explanations in the form of feature

importance scores to gain insights into the significance of different input features of the model’s results. As a baseline comparison, they implement a perplexity-based classifier that categorizes text based on its perplexity score, where GPT-2 is used for calculating perplexity scores. Their results show that the DistilBERT-based detector outperforms the perplexity-based classifier. However, its performance decreases when considering the rephrased dataset by ChatGPT.

In Liao et al. (2023), different models are proposed to detect medical text generated by ChatGPT: a fine-tuned BERT model (Devlin et al., 2019), a model based on Classification and Regression Trees (CART), an XGBoost model (Chen and Guestrin, 2016) and a perplexity classifier that utilizes BioGPT (Luo et al., 2022) for calculating text perplexity. Predictions by the BERT model are explained by visualizing the local features of the samples, where it can be seen that using conjuncts is an essential feature for the model classifying a medical text as machine-generated.

Liu et al. (2023) fine-tune RoBERTa to detect argumentative essays generated by different GPT models, including ChatGPT, and evaluate its performance on document, paragraph, and sentence-level classification. The essays are broken down into paragraphs and sentences for paragraph and sentence-level classification. They train and compare the performance of SVM models using different linguistic features. These models serve as a baseline to compare with the RoBERTa model and to understand which linguistic features differentiate between human and ChatGPT-generated text.

Guo et al. (2023) implement a machine learning and deep learning-based detector. They utilize a logistic regression model trained on the GLTR Test-2 dataset (Gehrmann et al., 2019) and two deep classifiers based on fine-tuning the pre-trained transformer model RoBERTa. One deep classifier is designed explicitly for single-text detection, while the other is intended for QA detection. The authors construct various training and testing datasets versions to assess the models’ robustness. They create full-text, sentence-level, and mixed subsets of the collected corpus. Each subset has both a raw version and a filtered version where prominent indicating words referring to humans (such as “Nope” and “Hmm”) or ChatGPT words (such as “AI assistant”) are removed. The evaluation of the models reveals that the RoBERTa-based models outper-

³<https://www.kaggle.com/datasets/chaitanyakck/medical-text>

Paper	Dataset	Approaches	Explainability	Code
Mitrović et al. 2023	Mitrović et al. 2023	DistilBERT PBC	SHAP	×
Liao et al. 2023	Liao et al. 2023	BERT PBC XGBoost CART	transformer-interpret	×
Liu et al. 2023	Liu et al. 2023 (ArguGPT)	RoBERTa-large SVM	×	✓*
Guo et al. 2023	Guo et al. 2023 (HC3)	GLTR RoBERTa-single RoBERTa-QA	×	✓
Antoun et al. 2023a	Antoun et al. 2023a Guo et al. 2023 (HC3)	CamemBERT CamemBERTa RoBERTa ELECTRA XLM-R	×	✓
Vasilatos et al. 2023	Ibrahim et al. 2023	PBC	×	×

Table 2: Methods proposed in the literature for detecting ChatGPT-generated text. PBC: Perplexity-based classifier. Publicly available models can be accessed by clicking on the ✓ character. *Authors indicate it will be made available at a future date.

form GLTR in terms of performance and exhibit more robustness against interference. Moreover, the RoBERTa-based models are not influenced by indicating words.

Building upon the work of Guo et al. (2023), Antoun et al. (2023a) propose an approach for developing robust detectors able to detect ChatGPT-generated text in different languages, with a focus on French. Their approach consists of fine-tuning pre-trained transformer-based models on English, French, and multilingual datasets. They train RoBERTa and ELECTRA (Clark et al., 2020) models on the English dataset, CamemBERT (Martin et al., 2020) and CamemBERTa (Antoun et al., 2023b) on the French datasets and XLM-R (Conneau et al., 2020) on the combined English and French dataset. They evaluate the robustness of these models against adversarial attacks, such as replacing characters with homoglyphs and adding misspelled words. Considering in-domain text, their results show that French models perform well in detecting machine-generated text. Still, they were outperformed by the English models, while XLM-R provides the best and most resilient performance against adversarial attacks for both

English and French. However, this performance decreases when evaluated on out-of-domain text.

Another method proposed for detecting ChatGPT-generated text is a metric-based approach proposed by Vasilatos et al. (2023) to detect machine-generated student assignments by calculating perplexity scores using GPT-2. They show that having category-wise thresholds (derived from dataset metadata) results in better detection performance than only having one threshold value.

3.3 Analysis of Human and ChatGPT-Generated Text

The textual characteristics of ChatGPT-generated text as well as its syntactic and linguistic features, are of significant focus in the works we reviewed. These linguistic and stylistic features are compared to the human-written texts in the datasets. In this section, we summarize and provide an overview of the findings of such analyses for the different domains and datasets we reviewed.

- **Medical domain:** Medical texts generated by ChatGPT have lower text perplexity and are more fluent, neutral, positive, and logical but more general in content and language style,

while medical texts written by humans are more diverse and specific (Liao et al., 2023).

- **English argumentative essays:** ChatGPT produces syntactically more complex sentences than English language learners, but ChatGPT-authored essays tend to have lower lexical diversity (Liu et al., 2023).
- **Multi-domain question answering:** ChatGPT writes in an organized and neutral way, offers less bias and harmful information, and refuses to answer questions where it believes it does not know. ChatGPT answers are formal, less emotional, and more objective than human answers (Guo et al., 2023).
- **Scientific abstracts:** ChatGPT has a better choice of vocabulary, can generate more unique words, uses more connecting words, and has fewer grammatical errors (Yu et al., 2023).
- **Language-agnostic characteristics:** The linguistic and syntactic characteristics of ChatGPT-generated text tend to be language-agnostic. Text generated in different languages, such as English, French, and Chinese, shows similar characteristics where ChatGPT tends to produce didactic and impersonal text without errors. Such errors can indicate human text, like grammatical, spelling or punctuation mistakes (Antoun et al., 2023a; Guo et al., 2023).

3.4 General Insights

Based on trends and regular mentions we encountered during the creation of our review, we now report some general insights on the state of detecting ChatGPT-generated text.

Role of explainable AI: Explainability techniques such as SHAP are helpful with detection models. These techniques provide insights into the most important features and words that contribute to classification, thus allowing a better understanding of the writing styles of humans and ChatGPT. This is also valuable in debugging detectors as they can highlight the main words contributing to the misclassification and thus enable better analysis of such models.

Humans versus ChatGPT in detection task: Another insight is that humans are worse at detecting machine-generated text by ChatGPT compared to ChatGPT itself. With additional training, humans would achieve better results.

Robustness of detectors: The robustness of detectors improves when they are trained on datasets that are extended to include also perturbed data, such as homoglyphs and misspellings. This might help the detectors focus more on writing style than writing errors. When evaluated on out-of-domain texts, the performance of detectors tends to decrease, especially when adversarial text is included.

Impact of text length on detection: The shorter the text length, the more challenging and less reliable detection becomes. Models trained on datasets containing full text and question-answer subsets (including answer contexts) do not perform well when evaluated on short texts such as sentences or smaller QA subsets.

Lack of special prompts in ChatGPT-generated text: Some conclusions and analyses in the reviewed papers have been made based on considering text generated by ChatGPT using its most general style and state, i.e., without using any special prompts that could ask ChatGPT to pretend to be a certain writer or to write in a special style. This could be an interesting area of investigation for future work, where new datasets are constructed, and the robustness of detectors against this type of text is tested.

Perplexity-based detectors: perplexity-based detectors depend on using open-source LLMs like GPT-2 and BioGPT to calculate perplexity scores. As ChatGPT generates the target text, calculating these scores using ChatGPT could benefit a lot in this task, as seen with other models using this method. However, this is not possible due to the unfortunate fact of it being a closed-source model.

Cost of constructing machine-generated datasets: Constructing and utilizing large-scale ChatGPT-generated datasets is important for drawing more generalized and precise conclusions. Therefore using ChatGPT's API is essential for this sake. However, the costs of doing so can be prohibitive.

Multilinguality: Our sample of papers has English dominance and performance for other languages is worse. Just as in NLP in general (Artetxe

et al., 2020), we call for more work in this area. This could help explain why some detectors are less reliable in detecting machine-generated text when the text is translated into different languages.

4 Conclusion and Future Work

The impressive capabilities of ChatGPT in producing high-quality and convincing text have brought attention to the risks associated with its improper usage across different domains. Consequently, the reliable detection of ChatGPT-generated text has become an important task. To address this concern, numerous datasets and detection methods have been proposed. In this paper, we provided a concise overview of the diverse datasets created, proposed methods, and qualitative insights of comparing human-written text with text generated by ChatGPT.

We see a wide variety of approaches and datasets in the papers we discussed. On the one hand, this is good to see since many factors, such as the domain, language, or format, influence the detection task. On the other hand, we also see a big diversity in experimental and dataset setups. Some works use adversarial examples, and others do not. Some allow the rephrasing of human text by ChatGPT, while others use purely human versus machine-generated text. Some works include the prompts and ChatGPT versions they used to generate the data; others do not. These, among other differences, make comparisons difficult, which is one reason we do not include scores in this survey. This also highlights important future work, namely to test *methods across datasets* and *datasets across methods*.

Another factor to consider is the domain of the text. The datasets we have discussed are in diverse domains and cover at least two important ones affected by ChatGPT’s risks: health and education. One notable domain we did not encounter is (fake) news. Although this is a big NLP field on its own, we expected more attention for it in the context of ChatGPT. Future work can definitely help in this area. The format of the text is related to the domain and is another important factor to consider. For example, the shared tasks we mentioned provide tweets, news articles, or reviews as their formats. A systematic look at format and domain influence concerning ChatGPT could be valuable future work.

Multilinguality is another open problem. As

with virtually all NLP tasks, we have seen that English is, unfortunately, the dominant language in the datasets. Experiments and gathering datasets across different languages are important future directions. The current task could also draw inspiration from the field of machine translation. It has a long and ongoing history of trying to detect (badly) translated text, so-called *translationese* (Baroni and Bernardini, 2006), which could be used or adapted to detect general machine-generated text.

Lastly, an important factor we have not seen discussed much is the temporal aspect of ChatGPT. Outputs might change over time, especially since it is a closed-source system. This calls for repeated tests over time to ensure detection methods are not regressing in their performance. Machine-generated text detection is also a cat-and-mouse game; since models are optimized to mimic human language, detection becomes harder and harder.

5 Limitations

A limitation of our work is that recent methods proposed for detecting ChatGPT-generated text are pre-prints published in arXiv, due to the rapid pace of work in this area. Additionally, we limit our scope to academic papers and exclude online non-academic tools as we do not know how those tools were trained or how they work internally.

This is also a big problem when discussing ChatGPT in general. Since it is a closed-sourced system without detailed information about its training and dataset, it is impossible to know if the results are reproducible. Models can change at any moment in the background, models can be decommissioned, or the price of access can change drastically. We are well aware of and concerned about these developments, but given the significant opportunities and risks ChatGPT poses, we believe a survey like this one is valuable.

Acknowledgements

We thank Florian Matthes and the Software Engineering for Business Information Systems (SEBIS) chair at TUM for their funding and support. We also thank the anonymous reviewers for their helpful and insightful comments.

References

Nash Anderson, Daniel L Belavy, Stephen M Perle, Sharief Hendricks, Luiz Hespanhol, Evert Verhagen, and Aamir R Memon. 2023. [AI did not write](#)

- this manuscript, or did it? Can we trick the AI text detector into generated texts? The potential future of ChatGPT and AI in Sports & Exercise Medicine manuscript generation. *BMJ Open Sport & Exercise Medicine*, 9(1):e001568.
- Wissam Antoun, Virginie Mouilleron, Benoît Sagot, and Djamé Seddah. 2023a. [Towards a Robust Detection of Language Model Generated Text: Is ChatGPT that Easy to Detect?](#) (arXiv:2306.05871v1).
- Wissam Antoun, Benoît Sagot, and Djamé Seddah. 2023b. [Data-Efficient French Language Modeling with CamemBERTa](#). (arXiv:2306.01497v1).
- Mikel Artetxe, Sebastian Ruder, Dani Yogatama, Gorka Labaka, and Eneko Agirre. 2020. [A Call for More Rigor in Unsupervised Cross-lingual Learning](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7375–7388, Online. Association for Computational Linguistics.
- Marco Baroni and Silvia Bernardini. 2006. [A New Approach to the Study of Translationese: Machine-learning the Difference between Original and Translated Text](#). *Literary and Linguistic Computing*, 21(3):259–274.
- Daniel Blanchard, Joel Tetreault, Derrick Higgins, Aoife Cahill, and Martin Chodorow. 2013. [TOEFL11: A corpus of non-native english](#). *ETS Research Report Series*, 2013(2):i–15.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D 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 Ziegler, Jeffrey Wu, Clemens Winter, Chris 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 *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Tianqi Chen and Carlos Guestrin. 2016. [XGBoost: A Scalable Tree Boosting System](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 785–794, New York, NY, USA. Association for Computing Machinery.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pilla, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [PaLM: Scaling Language Modeling with Pathways](#). (arXiv:2204.02311v5).
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: Pre-training Text Encoders as Discriminators Rather Than Generators](#). (arXiv:2003.10555v1).
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised Cross-lingual Representation Learning at Scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Debby R. E. Cotton, Peter A. Cotton, and J. Reuben Shipway. 2023. [Chatting and cheating: Ensuring academic integrity in the era of ChatGPT](#). *Innovations in Education and Teaching International*, 0(0):1–12.
- Evan Crothers, Nathalie Japkowicz, and Herna Viktor. 2023. [Machine Generated Text: A Comprehensive Survey of Threat Models and Detection Methods](#). (arXiv:2210.07321v4).
- Maxime De Bruyn, Ehsan Lotfi, Jeska Buhmann, and Walter Daelemans. 2021. [MFAQ: a multilingual FAQ dataset](#). In *Proceedings of the 3rd Workshop on Machine Reading for Question Answering*, pages 1–13, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). (arXiv:1810.04805v2).
- Liam Dugan, Daphne Ippolito, Arun Kirubakaran, and Chris Callison-Burch. 2020. [RoFT: A Tool for Evaluating Human Detection of Machine-Generated Text](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 189–196, Online. Association for Computational Linguistics.
- Liam Dugan, Daphne Ippolito, Arun Kirubakaran, Sherry Shi, and Chris Callison-Burch. 2023. [Real or Fake Text?: Investigating Human Ability to Detect Boundaries between Human-Written and Machine-Generated Text](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(11):12763–12771.

- Angela Fan, Yacine Jernite, Ethan Perez, David Grangier, Jason Weston, and Michael Auli. 2019. [ELI5: Long Form Question Answering](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3558–3567, Florence, Italy. Association for Computational Linguistics.
- Catherine A. Gao, Frederick M. Howard, Nikolay S. Markov, Emma C. Dyer, Siddhi Ramesh, Yuan Luo, and Alexander T. Pearson. 2022. [Comparing scientific abstracts generated by ChatGPT to original abstracts using an artificial intelligence output detector, plagiarism detector, and blinded human reviewers](#). page 2022.12.23.521610.
- 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.
- Biyang Guo, Xin Zhang, Ziyuan Wang, Minqi Jiang, Jinran Nie, Yuxuan Ding, Jianwei Yue, and Yupeng Wu. 2023. [How Close is ChatGPT to Human Experts? Comparison Corpus, Evaluation, and Detection](#). (arXiv:2301.07597v1).
- Xinlei He, Xinyue Shen, Zeyuan Chen, Michael Backes, and Yang Zhang. 2023. [MGTBench: Benchmarking Machine-Generated Text Detection](#). (arXiv:2303.14822v2).
- Hazem Ibrahim, Fengyuan Liu, Rohail Asim, Balaraju Battu, Sidahmed Benabderrahmane, Bashar Alhafni, Wifag Adnan, Tuka Alhanai, Bedoor AlShibli, Riyadh Baghdadi, Jocelyn J. Bélanger, Elena Beretta, Kemal Celik, Moumena Chaqfeh, Mohammed F. Daqaq, Zaynab El Bernoussi, Daryl Fougner, Borja Garcia de Soto, Alberto Gandolfi, Andras Gyorgy, Nizar Habash, J. Andrew Harris, Aaron Kaufman, Lefteris Kirousis, Korhan Kocak, Kangsan Lee, Seungah S. Lee, Samreen Malik, Michail Maniatakos, David Melcher, Azzam Mourad, Minsu Park, Mahmoud Rasras, Alicja Reuben, Dania Zantout, Nancy W. Gleason, Kinga Makovi, Talal Rahwan, and Yasir Zaki. 2023. [Perception, performance, and detectability of conversational artificial intelligence across 32 university courses](#). (arXiv:2305.13934v1).
- 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.
- Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Liwei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. [MIMIC-III, a freely accessible critical care database](#). *Scientific Data*, 3(1):160035.
- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. 2023a. [A Watermark for Large Language Models](#). (arXiv:2301.10226v3).
- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Manli Shu, Khalid Saifullah, Kezhi Kong, Kasun Fernando, Aniruddha Saha, Micah Goldblum, and Tom Goldstein. 2023b. [On the Reliability of Watermarks for Large Language Models](#). (arXiv:2306.04634v3).
- Tomáš Kočický, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. [The NarrativeQA Reading Comprehension Challenge](#). *Transactions of the Association for Computational Linguistics*, 6:317–328.
- Sachin Kumar, Vidhisha Balachandran, Lucille Njoo, Antonios Anastasopoulos, and Yulia Tsvetkov. 2023. [Language Generation Models Can Cause Harm: So What Can We Do About It? An Actionable Survey](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3299–3321, Dubrovnik, Croatia. Association for Computational Linguistics.
- Wenxiong Liao, Zhengliang Liu, Haixing Dai, Shaochen Xu, Zihao Wu, Yiyang Zhang, Xiaoke Huang, Dajiang Zhu, Hongmin Cai, Tianming Liu, and Xiang Li. 2023. [Differentiate ChatGPT-generated and Human-written Medical Texts](#). (arXiv:2304.11567v1).
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [TruthfulQA: Measuring How Models Mimic Human Falsehoods](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.
- Yikang Liu, Ziyin Zhang, Wanyang Zhang, Shisen Yue, Xiaojing Zhao, Xinyuan Cheng, Yiwen Zhang, and Hai Hu. 2023. [ArguGPT: Evaluating, understanding and identifying argumentative essays generated by GPT models](#). (arXiv:2304.07666v1).
- Scott M Lundberg and Su-In Lee. 2017. [A Unified Approach to Interpreting Model Predictions](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Renqian Luo, Liai Sun, Yingce Xia, Tao Qin, Sheng Zhang, Hoifung Poon, and Tie-Yan Liu. 2022. [BioGPT: Generative Pre-trained Transformer for Biomedical Text Generation and Mining](#). *Briefings in Bioinformatics*, 23(6):bbac409.

- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamel Seddah, and Benoît Sagot. 2020. [CamemBERT: A Tasty French Language Model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. 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](#). (arXiv:2301.11305v2).
- Sandra Mitrović, Davide Andreoletti, and Omran Ayoub. 2023. [ChatGPT or Human? Detect and Explain](#). Explaining Decisions of Machine Learning Model for Detecting Short ChatGPT-generated Text. (arXiv:2301.13852v1).
- Arpit Narechania, Alireza Karduni, Ryan Wesslen, and Emily Wall. 2022. [VitaLITY: Promoting Serendipitous Discovery of Academic Literature with Transformers & Visual Analytics](#). *IEEE Transactions on Visualization and Computer Graphics*, 28(1):486–496.
- OpenAI. 2022. [ChatGPT: Optimizing Language Models for Dialogue](#).
- Alessandro Pegoraro, Kavita Kumari, Hossein Fereidooni, and Ahmad-Reza Sadeghi. 2023. [To ChatGPT, or not to ChatGPT: That is the question!](#) (arXiv:2304.01487v2).
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2018. [Language Models are Unsupervised Multitask Learners](#).
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100,000+ Questions for Machine Comprehension of Text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [DistilBERT, a distilled version of BERT: Smaller, faster, cheaper and lighter](#). (arXiv:1910.01108v4).
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2021. [Societal Biases in Language Generation: Progress and Challenges](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4275–4293, Online. Association for Computational Linguistics.
- 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](#). (arXiv:1908.09203v2).
- Yi Tay, Dara Bahri, Che Zheng, Clifford Brunk, Donald Metzler, and Andrew Tomkins. 2020. [Reverse Engineering Configurations of Neural Text Generation Models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 275–279, Online. Association for Computational Linguistics.
- 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.
- Christoforos Vasilatos, Manaar Alam, Talal Rahwan, Yasir Zaki, and Michail Maniatakos. 2023. [HowkGPT: Investigating the Detection of ChatGPT-generated University Student Homework through Context-Aware Perplexity Analysis](#). (arXiv:2305.18226v2).
- Laura Weidinger, Jonathan Uesato, Maribeth Rauh, Conor Griffin, Po-Sen Huang, John Mellor, Amelia Glaese, Myra Cheng, Borja Balle, Atoosa Kasirzadeh, Courtney Biles, Sasha Brown, Zac Kenton, Will Hawkins, Tom Stepleton, Abeba Birhane, Lisa Anne Hendricks, Laura Rimell, William Isaac, Julia Haas, Sean Legassick, Geoffrey Irving, and Iason Gabriel. 2022. [Taxonomy of Risks posed by Language Models](#). In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT '22*, pages 214–229, New York, NY, USA. Association for Computing Machinery.
- Luoxuan Weng, Minfeng Zhu, Kam Kwai Wong, Shi Liu, Jiashun Sun, Hang Zhu, Dongming Han, and Wei Chen. 2023. [Towards an Understanding and Explanation for Mixed-Initiative Artificial Scientific Text Detection](#). (arXiv:2304.05011v1).
- Yi Yang, Wen-tau Yih, and Christopher Meek. 2015. [WikiQA: A Challenge Dataset for Open-Domain Question Answering](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2013–2018, Lisbon, Portugal. Association for Computational Linguistics.
- Peipeng Yu, Jiahua Chen, Xuan Feng, and Zhihua Xia. 2023. [CHEAT: A Large-scale Dataset for Detecting ChatGPT-writtEn AbsTracts](#). (arXiv:2304.12008v1).
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. 2020. [Defending Against Neural Fake News](#). (arXiv:1905.12616v3).
- Xu Zhi-jia. 2008. [A Review of Spoken and Written English Corpus of Chinese Learners](#). *Journal of Longyan University*.

Terry Yue Zhuo, Yujin Huang, Chunyang Chen, and Zhenchang Xing. 2023. [Red teaming ChatGPT via Jailbreaking: Bias, Robustness, Reliability and Toxicity](#). (arXiv:2301.12867v4).