

GPT detectors are biased against non-native English writers

Weixin Liang^{1*}, Mert Yuksekgonul^{1*}, Yining Mao^{2*}, Eric Wu^{2*}, and James Zou^{1,2,3,+}

¹Department of Computer Science, Stanford University, Stanford, CA, USA

²Department of Electrical Engineering, Stanford University, Stanford, CA, USA

³Department of Biomedical Data Science, Stanford University, Stanford, CA, USA

⁺Correspondence should be addressed to: jamesz@stanford.edu

^{*}these authors contributed equally to this work

ABSTRACT

The rapid adoption of generative language models has brought about substantial advancements in digital communication, while simultaneously raising concerns regarding the potential misuse of AI-generated content. Although numerous detection methods have been proposed to differentiate between AI and human-generated content, the fairness and robustness of these detectors remain underexplored. In this study, we evaluate the performance of several widely-used GPT detectors using writing samples from native and non-native English writers. Our findings reveal that these detectors consistently misclassify non-native English writing samples as AI-generated, whereas native writing samples are accurately identified. Furthermore, we demonstrate that simple prompting strategies can not only mitigate this bias but also effectively bypass GPT detectors, suggesting that GPT detectors may unintentionally penalize writers with constrained linguistic expressions. Our results call for a broader conversation about the ethical implications of deploying ChatGPT content detectors and caution against their use in evaluative or educational settings, particularly when they may inadvertently penalize or exclude non-native English speakers from the global discourse. The published version of this study can be accessed at: [www.cell.com/patterns/fulltext/S2666-3899\(23\)00130-7](http://www.cell.com/patterns/fulltext/S2666-3899(23)00130-7)

Introduction

Generative language models based on GPT, such as ChatGPT¹, have taken the world by storm. Within a mere two months of its launch, ChatGPT attracted over 100 million monthly active users, making it one of the fastest-growing consumer internet applications in history^{2,3}. While these powerful models offer immense potential for enhancing productivity and creativity⁴⁻⁶, they also introduce the risk of AI-generated content being passed off as human-written, which may lead to potential harms, such as the spread of fake content and exam cheating⁷⁻¹¹.

Recent studies reveal the challenges humans face in detecting AI-generated content, emphasizing the urgent need for effective detection methods^{7-9,12}. Although several publicly available GPT detectors have been developed to mitigate the risks associated with AI-generated content, their effectiveness and reliability remain uncertain due to limited evaluation¹³⁻²¹. This lack of understanding is particularly concerning given the potentially damaging consequences of misidentifying human-written content as AI-generated, especially in educational settings^{22,23}.

Given the transformative impact of generative language models and the potential risks associated with their misuse, developing trustworthy and accurate detection methods is crucial. In this study, we evaluate several publicly available GPT detectors on writing samples from native and non-native English writers. We uncover a concerning pattern: GPT detectors consistently misclassify non-native English writing samples as AI-generated while not making the same mistakes for native writing samples. Further investigation reveals that simply prompting GPT to generate more linguistically diverse versions of the non-native samples effectively removes this bias, suggesting that GPT detectors may inadvertently penalize writers with limited linguistic expressions.

Our findings emphasize the need for increased focus on the fairness and robustness of GPT detectors, as overlooking their biases may lead to unintended consequences, such as the marginalization of non-native speakers in evaluative or educational settings. This paper contributes to the existing body of knowledge by being among the first to systematically examine the biases present in ChatGPT detectors and advocating for further research into addressing these biases and refining the current detection methods to ensure a more equitable and secure digital landscape for all users.

Results

GPT detectors exhibit bias against non-native English authors

We evaluated the performance of seven widely-used GPT detectors on a corpus of 91 human-authored TOEFL essays obtained from a Chinese educational forum and 88 US 8-th grade essays sourced from the Hewlett Foundation's Automated Student Assessment Prize (ASAP) dataset²⁴ (Fig. 1a). The detectors demonstrated near-perfect accuracy for US 8-th grade essays. However, they misclassified over half of the TOEFL essays as "AI-generated" (average false positive rate: 61.22%). All seven detectors unanimously identified 18 of the 91 TOEFL essays (19.78%) as AI-authored, while 89 of the 91 TOEFL essays (97.80%) are flagged as AI-generated by at least one detector. For the TOEFL essays that were unanimously identified (Fig. 1b), we observed that they had significantly lower perplexity compared to the others (P-value: 9.74E-05). This suggests that GPT detectors may penalize non-native writers with limited linguistic expressions.

Mitigating Bias through Linguistic Diversity Enhancement of Non-Native Samples

To explore the hypothesis that the restricted linguistic variability and word choices characteristic of non-native English writers contribute to the observed bias, we employed ChatGPT to enrich the language in the TOEFL essays, aiming to emulate the vocabulary usage of native speakers (Prompt: "Enhance the word choices to sound more like that of a native speaker.") (Fig. 1c). Remarkably, this intervention led to a substantial reduction in misclassification, with the average false positive rate decreasing by 49.45% (from 61.22% to 11.77%). Post-intervention, the TOEFL essays' perplexity significantly increased (P-value=9.36E-05), and only 1 out of 91 essays (1.10%) was unanimously detected as AI-written. In contrast, applying ChatGPT to adjust the word choices in US 8th-grade essays to mimic non-native speaker writing (Prompt: "Simplify word choices as if written by a non-native speaker.") led to a significant increase in the misclassification rate as AI-generated text, from an average of 5.19% across detectors to 56.65% (Fig. 1ac). This word choice adjustment also resulted in significantly lower text perplexity (Fig. 1d).

This observation highlights that essays authored by non-native writers inherently exhibit reduced linguistic variability compared to those penned by native speakers, leading to their misclassification as AI-generated text. Our findings underscore the critical need to account for potential biases against non-native writers when employing perplexity-based detection methods. Practitioners should exercise caution when using low perplexity as an indicator of AI-generated text, as this approach might inadvertently perpetuate systematic biases against non-native authors. Non-native English writers have been shown to exhibit reduced linguistic variability in terms of lexical richness²⁵, lexical diversity^{26,27}, syntactic complexity²⁸⁻³⁰, and grammatical complexity³¹. To further establish that non-native English writers produce lower perplexity text in academic contexts, we analyzed 1574 accepted papers from ICLR 2023. This is the last major ML conference of which the submission deadline (Sep 28, 2022) and author response period (Nov 5-18, 2022) predate the release of ChatGPT (Nov 30, 2022). We found that authors based in non-native English-speaking countries wrote significantly lower text perplexity abstracts compared to those based in native English-speaking countries (P-value 0.035). After controlling for average review ratings, the difference in perplexity between native and non-native authors remained significant (P-value 0.033). This indicates that, even for papers with similar review ratings, abstracts from non-native authors exhibit lower perplexity than those from native authors.

Simple prompt can easily bypass current GPT detectors

Enhancing linguistic diversity can help to not only mitigate the bias for non-native English writers, but also make GPT-generated content bypass GPT detectors. As a proof of concept, we prompted ChatGPT-3.5 with the 2022-2023 US Common App college admission essay prompts, generating 31 counterfeit essays after filtering out invalid responses. While detectors were initially effective, a second-round self-edit prompt ("Elevate the provided text by employing literary language") applied to ChatGPT-3.5 significantly reduced detection rates from 100% to 13% (Fig. 2a). Although ChatGPT-3.5 generated essays initially exhibit notably low perplexity, applying the self-edit prompt leads to a significant increase in perplexity (Fig. 2b) (P-value 1.94E-15). In a parallel experiment, we prompted ChatGPT-3.5 to generate scientific abstracts using 145 Stanford CS224n final project report titles (Fig. 2c). Detectors were less effective in this context, partly because the generated abstracts have slightly higher perplexity than their essays counterpart (Figs. 2bd), but still identified up to 68% of fake abstracts. However, applying a second-round self-edit prompt ("Elevate the provided text by employing advanced technical language") lowered detection rates to up to 28%. Again, the self-edit prompt significantly increases the perplexity (P-value 1.06E-31). These results demonstrate the perplexity of GPT-generated text can be significantly improved using straightforward prompt design, and thus easily bypass current GPT detectors, revealing the vulnerability of perplexity-based approaches. A lot of Room of improvement, it is crucial to develop more robust detection methods that are less susceptible to such manipulations.

Discussion

This study reveals a notable bias in GPT detectors against non-native English writers, as evidenced by the high misclassification rate of non-native-authored TOEFL essays, in stark contrast to the near zero misclassification rate of college essays, which are

presumably authored by native speakers. One possible explanation of this discrepancy is that non-native authors exhibited limited linguistic variability and word choices, which consequently result in lower perplexity text. Non-native English writers have been shown to exhibit reduced linguistic variability in terms of lexical richness²⁵, lexical diversity^{26,27}, syntactic complexity^{28–30}, and grammatical complexity³¹. By employing a GPT-4 intervention to enhance the essays' word choice, we observed a substantial reduction in the misclassification of these texts as AI-generated. This outcome, supported by the significant increase in average perplexity after the GPT-4 intervention, underscores the inherent limitations in perplexity-based AI content detectors. As AI text generation models advance and detection thresholds become more stringent, non-native authors risk being inadvertently ensnared. Paradoxically, to evade false detection as AI-generated content, these writers may need to rely on AI tools to refine their vocabulary and linguistic diversity. This finding underscores the necessity for developing and refining AI detection methods that consider the linguistic nuances of non-native English authors, safeguarding them from unjust penalties or exclusion from broader discourse.

Our investigation into the effectiveness of simple prompts in bypassing GPT detectors, along with recent studies on paraphrasing attacks^{32,33}, raises significant concerns about the reliability of current detection methods. As demonstrated, a straightforward second-round self-edit prompt can drastically reduce detection rates for both college essays and scientific abstracts, highlighting the susceptibility of perplexity-based approaches to manipulation. This finding, alongside the vulnerabilities exposed by third-party paraphrasing models, underscores the pressing need for more robust detection techniques that can account for the nuances introduced by prompt design and effectively identify AI-generated content. Ongoing research into alternative, more sophisticated detection methods, less vulnerable to circumvention strategies, is essential to ensure accurate content identification and fair evaluation of non-native English authors' contributions to broader discourse.

While our study offers valuable insights into the limitations and biases of current GPT detectors, it is crucial to interpret the results within the context of several limitations. Firstly, although our datasets and analysis present novel perspectives as a pilot study, the sample sizes employed in this research are relatively small. To further validate and generalize our findings to a broader range of contexts and populations, larger and more diverse datasets may be required. Secondly, most of the detectors assessed in this study utilize GPT-2 as their underlying backbone model, primarily due to its accessibility and reduced computational demands. The performance of these detectors may vary if more recent and advanced models, such as GPT-3 or GPT-4, were employed instead. Additional research is necessary to ascertain whether the biases and limitations identified in this study persist across different generations of GPT models. Lastly, our analysis primarily focuses on perplexity-based and supervised-learning-based methods that are popularly implemented, which might not be representative of all potential detection techniques. For instance, DetectGPT¹⁷, based on second-order log probability, has exhibited improved performance in specific tasks but is orders of magnitude more computationally demanding to execute, and thus not widely deployed at scale. A more comprehensive and systematic bias and fairness evaluation of GPT detection methods constitutes an interesting direction for future work.

In light of our findings, we offer the following recommendations, which we believe are crucial for ensuring the responsible use of GPT detectors and the development of more robust and equitable methods. First, we strongly caution against the use of GPT detectors in evaluative or educational settings, particularly when assessing the work of non-native English speakers. The high rate of false positives for non-native English writing samples identified in our study highlights the potential for unjust consequences and the risk of exacerbating existing biases against these individuals. Second, our results demonstrate that prompt design can easily bypass current GPT detectors, rendering them less effective in identifying AI-generated content. Consequently, future detection methods should move beyond solely relying on perplexity measures and consider more advanced techniques, such as second-order perplexity methods¹⁷ and watermarking techniques^{34,35}. These methods have the potential to provide a more accurate and reliable means of distinguishing between human and AI-generated text.

Correspondence

Correspondence should be addressed to J.Z. (email: jamesz@stanford.edu).

Competing interests

The authors declare no conflict of interest.

Acknowledgements

We thank B. He for discussions. J.Z. is supported by the National Science Foundation (CCF 1763191 and CAREER 1942926), the US National Institutes of Health (P30AG059307 and U01MH098953) and grants from the Silicon Valley Foundation and the Chan-Zuckerberg Initiative.

References

1. OpenAI. ChatGPT. <https://chat.openai.com/> (2022). Accessed: 2022-12-31.
2. Hu, K. Chatgpt sets record for fastest-growing user base - analyst note. *Reuters* (2023).
3. Paris, M. Chatgpt hits 100 million users, google invests in ai bot and catgpt goes viral. *Forbes* (2023).
4. Lee, M. *et al.* Evaluating human-language model interaction. *arXiv preprint arXiv:2212.09746* (2022).
5. Kung, T. H. *et al.* Performance of chatgpt on usmle: Potential for ai-assisted medical education using large language models. *PLoS digital health* **2**, e0000198 (2023).
6. Terwiesch, C. Would chat gpt3 get a wharton mba? a prediction based on its performance in the operations management course. *Mack Inst. for Innov. Manag. at Whart. Sch. Univ. Pennsylvania* (2023).
7. Else, H. Abstracts written by chatgpt fool scientists. *Nature* (2023).
8. Gao, C. A. *et al.* Comparing scientific abstracts generated by chatgpt to original abstracts using an artificial intelligence output detector, plagiarism detector, and blinded human reviewers. *bioRxiv* 2022–12 (2022).
9. Kreps, S., McCain, R. & Brundage, M. All the news that's fit to fabricate: Ai-generated text as a tool of media misinformation. *J. Exp. Polit. Sci.* **9**, 104–117, DOI: [10.1017/XPS.2020.37](https://doi.org/10.1017/XPS.2020.37) (2022).
10. Editorial, N. Tools such as chatgpt threaten transparent science; here are our ground rules for their use. *Nature* **613**, 612–612 (2023).
11. ICML. Clarification on large language model policy LLM. <https://icml.cc/Conferences/2023/llm-policy> (2023).
12. Clark, E. *et al.* All that's 'human' is not gold: Evaluating human evaluation of generated text. 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)*, 7282–7296 (2021).
13. OpenAI. GPT-2: 1.5B release. <https://openai.com/research/gpt-2-1-5b-release> (2019). Accessed: 2019-11-05.
14. Jawahar, G., Abdul-Mageed, M. & Lakshmanan, L. V. Automatic detection of machine generated text: A critical survey. *arXiv preprint arXiv:2011.01314* (2020).
15. Fagni, T., Falchi, F., Gambini, M., Martella, A. & Tesconi, M. Tweepfake: About detecting deepfake tweets. *Plos one* **16**, e0251415 (2021).
16. Ippolito, D., Duckworth, D., Callison-Burch, C. & Eck, D. Automatic detection of generated text is easiest when humans are fooled. *arXiv preprint arXiv:1911.00650* (2019).
17. Mitchell, E., Lee, Y., Khazatsky, A., Manning, C. D. & Finn, C. DetectGPT: Zero-shot machine-generated text detection using probability curvature. *arXiv preprint arXiv:2301.11305* (2023).
18. Solaiman, I. *et al.* Release strategies and the social impacts of language models. *arXiv preprint arXiv:1908.09203* (2019).
19. Gehrmann, S., Strobel, H. & Rush, A. M. Gltr: Statistical detection and visualization of generated text. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 111–116 (2019).
20. Heikkilä, M. How to spot ai-generated text. *MIT Technol. Rev.* (2022).
21. Crothers, E., Japkowicz, N. & Viktor, H. Machine generated text: A comprehensive survey of threat models and detection methods. *arXiv preprint arXiv:2210.07321* (2022).
22. Rosenblatt, K. Chatgpt banned from new york city public schools' devices and networks. *NBC News* (2023). Accessed: 22.01.2023.
23. Kasneci, E. *et al.* Chatgpt for good? on opportunities and challenges of large language models for education. *Learn. Individ. Differ.* **103**, 102274 (2023).
24. Kaggle. The hewlett foundation: Automated essay scoring. <https://www.kaggle.com/c/asap-aes> (2012). Accessed: 2023-03-15.
25. Laufer, B. & Nation, P. Vocabulary size and use: Lexical richness in l2 written production. *Appl. linguistics* **16**, 307–322 (1995).
26. Jarvis, S. Short texts, best-fitting curves and new measures of lexical diversity. *Lang. Test.* **19**, 57–84 (2002).
27. Daller, H., Van Hout, R. & Treffers-Daller, J. Lexical richness in the spontaneous speech of bilinguals. *Appl. linguistics* **24**, 197–222 (2003).

28. Lu, X. A corpus-based evaluation of syntactic complexity measures as indices of college-level esl writers' language development. *TESOL quarterly* **45**, 36–62 (2011).
29. Crossley, S. A. & McNamara, D. S. Does writing development equal writing quality? a computational investigation of syntactic complexity in l2 learners. *J. Second. Lang. Writ.* **26**, 66–79 (2014).
30. Ortega, L. Syntactic complexity measures and their relationship to l2 proficiency: A research synthesis of college-level l2 writing. *Appl. linguistics* **24**, 492–518 (2003).
31. Biber, D., Gray, B. & Poonpon, K. Should we use characteristics of conversation to measure grammatical complexity in l2 writing development? *Tesol Q.* **45**, 5–35 (2011).
32. Krishna, K., Song, Y., Karpinska, M., Wieting, J. & Iyyer, M. Paraphrasing evades detectors of ai-generated text, but retrieval is an effective defense. *arXiv preprint arXiv:2303.13408* (2023).
33. Sadasivan, V. S., Kumar, A., Balasubramanian, S., Wang, W. & Feizi, S. Can ai-generated text be reliably detected? *arXiv preprint arXiv:2303.11156* (2023).
34. Kirchenbauer, J. *et al.* A watermark for large language models. *arXiv preprint arXiv:2301.10226* (2023).
35. Gu, C., Huang, C., Zheng, X., Chang, K.-W. & Hsieh, C.-J. Watermarking pre-trained language models with backdooring. *arXiv preprint arXiv:2210.07543* (2022).
36. Liang, W., Yuksekgonul, M., Mao, Y., Wu, E. & Zou, J. ChatGPT-Detector-Bias: v1.0.0, DOI: [10.5281/zenodo.7893958](https://doi.org/10.5281/zenodo.7893958) (2023).

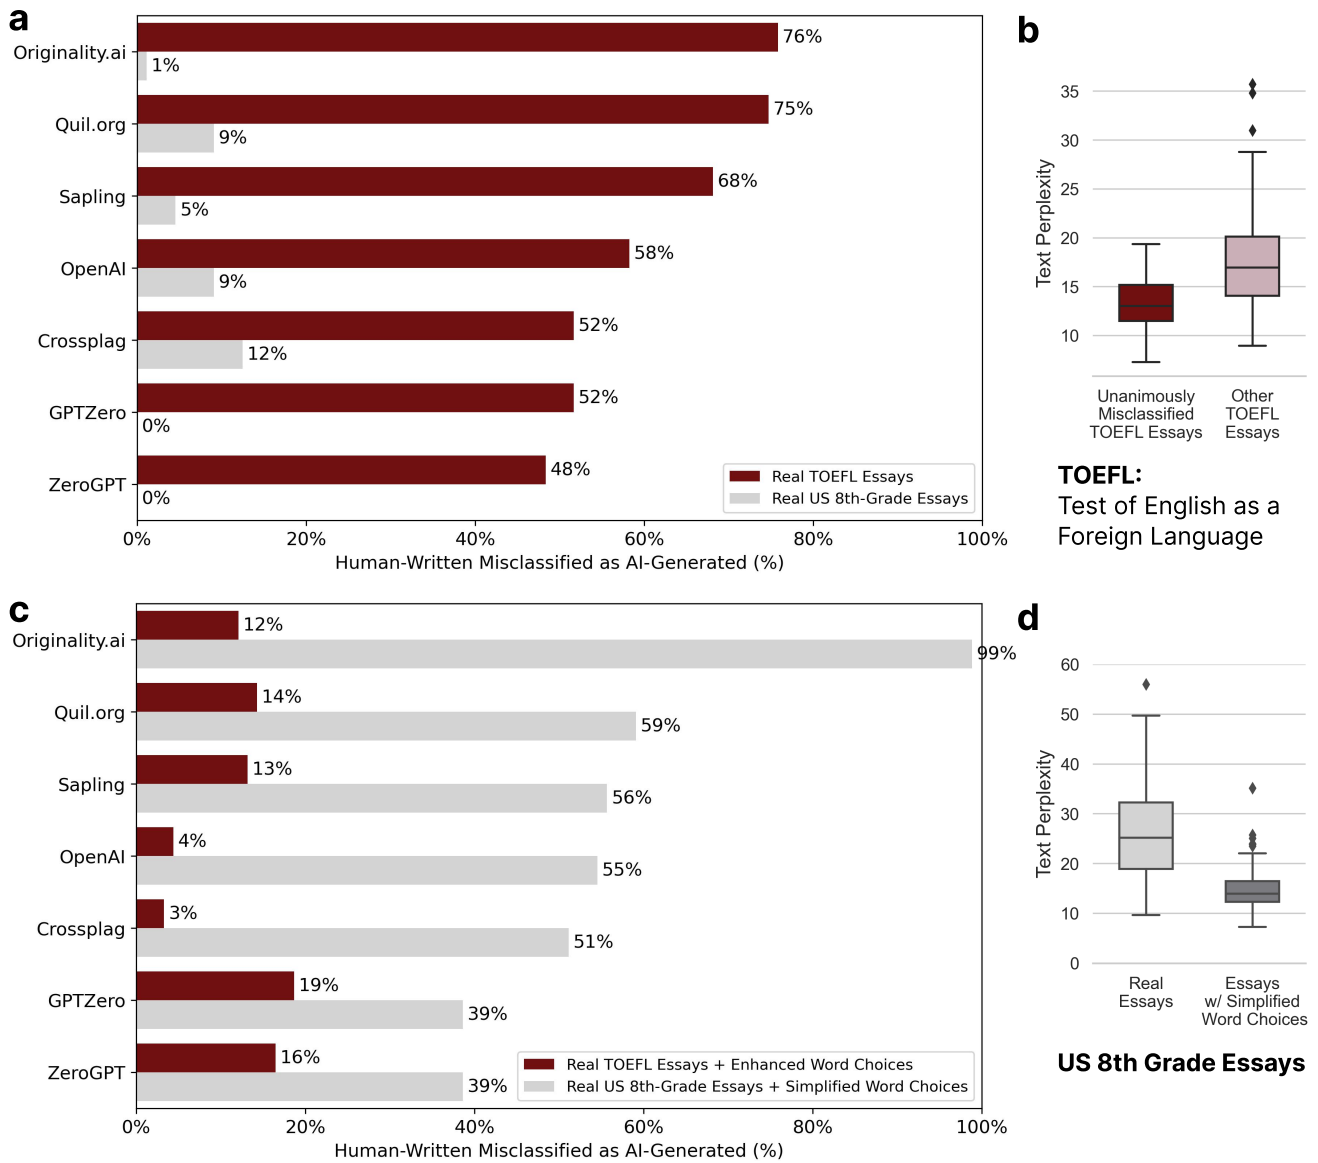


Figure 1. Bias in GPT detectors against non-native English writing samples. (a) Performance comparison of seven widely-used GPT detectors. More than half of the non-native-authored TOEFL (Test of English as a Foreign Language) essays are incorrectly classified as "AI-generated," while detectors exhibit near-perfect accuracy for US 8-th grade essays. (b) TOEFL essays unanimously misclassified as AI-generated show significantly lower perplexity compared to others, suggesting that GPT detectors might penalize authors with limited linguistic expressions. (c) Using ChatGPT to improve the word choices in TOEFL essays (Prompt: "Enhance the word choices to sound more like that of a native speaker.") significantly reduces misclassification as AI-generated text. Conversely, applying ChatGPT to simplify the word choices in US 8th-grade essays (Prompt: "Simplify word choices as if written by a non-native speaker.") significantly increases misclassification as AI-generated text. (d) The US 8th-grade essays with simplified word choices demonstrate significantly lower text perplexity.




Figure 2. Simple prompts effectively bypass GPT detectors. (a) For ChatGPT-3.5 generated college admission essays, the performance of seven widely-used GPT detectors declines markedly when a second-round self-edit prompt (“*Elevate the provided text by employing literary language*”) is applied, with detection rates dropping from up to 100% to up to 13%. (b) ChatGPT-3.5 generated essays initially exhibit notably low perplexity; however, applying the self-edit prompt leads to a significant increase in perplexity. (c) Similarly, in detecting ChatGPT-3.5 generated scientific abstracts, a second-round self-edit prompt (“*Elevate the provided text by employing advanced technical language*”) leads to a reduction in detection rates from up to 68% to up to 28%. (d) ChatGPT-3.5 generated abstracts have slightly higher perplexity than the generated essays but remain low. Again, the self-edit prompt significantly increases the perplexity.

Materials and Methods

Data availability

Our data, results, and code are available on both GitHub (<https://github.com/Weixin-Liang/ChatGPT-Detector-Bias/>) and Zenodo³⁶.

Evaluation of off-the-shelf GPT detectors

We assessed seven widely-used off-the-shelf GPT detectors:

1. Originality.AI: <https://app.originality.ai/api-access>
2. Quil.org: <https://aiwritingcheck.org/>
3. Sapling: <https://sapling.ai/ai-content-detector>
4. OpenAI: <https://openai-openai-detector.hf.space/>
5. Crossplag: <https://crossplag.com/ai-content-detector/>
6. GPTZero: <https://gptzero.me/>
7. ZeroGPT: <https://www.zerogpt.com/>

Accessed on March 15, 2023.

ChatGPT prompts used

1. **ChatGPT prompt for refining real TOEFL essays:** “Enhance the word choices to sound more like that of a native speaker: <TOEFL essay text>”
2. **ChatGPT prompt for adjusting real US 8th grade essays:** “Simplify word choices as of written by a non-native speaker.”
3. **ChatGPT prompts for the US college admission essays**

(a) **[1st round] ChatGPT prompt for generating US college admission essays:** “Hi GPT, I’d like you to write a college application essay. <college-essay-prompt>” where the <college-essay-prompt> corresponds to one of the Common App 2022-2023 essay prompts as follows (7 prompts in total):

- i. Some students have a background, identity, interest, or talent that is so meaningful they believe their application would be incomplete without it. If this sounds like you, then please share your story.
- ii. The lessons we take from obstacles we encounter can be fundamental to later success. Recount a time when you faced a challenge, setback, or failure. How did it affect you, and what did you learn from the experience?
- iii. Reflect on a time when you questioned or challenged a belief or idea. What prompted your thinking? What was the outcome?
- iv. Reflect on something that someone has done for you that has made you happy or thankful in a surprising way. How has this gratitude affected or motivated you?
- v. Discuss an accomplishment, event, or realization that sparked a period of personal growth and a new understanding of yourself or others.
- vi. Describe a topic, idea, or concept you find so engaging that it makes you lose all track of time. Why does it captivate you? What or who do you turn to when you want to learn more?
- vii. Share an essay on any topic of your choice. It can be one you’ve already written, one that responds to a different prompt, or one of your own design.

For each college essay prompt, we run 10 trials, resulting in 70 trials in total. After filtering out invalid responses (E.g., “As an AI language model, I don’t have a personal background, identity, interest or talent. Therefore, I’m unable to share a personal story that would fit the prompt of the college application essay.”), we obtained 31 counterfeit essays.

(b) **[2nd round] ChatGPT prompt for refining ChatGPT-generated US college admission essays:** “Elevate the provided text by employing literary language: <generated essay>” where the <generated essay> originates from the first round.

4. ChatGPT prompts for scientific abstracts

- (a) **[1st round] ChatGPT prompt for generating US college admission essays:** “Please draft an abstract (about 120 words) for a final report based on the title '<title>'” where the <title> is a scientific project title.
- (b) **[2nd round] ChatGPT prompt for refining ChatGPT-generated scientific abstracts:** “Elevate the provided text by employing advanced technical language: <generated abstract>” where the <generated abstract> comes from the first round.

We utilized the March 14 version of ChatGPT 3.5.

Data

TOEFL Essays

We collected a total of 91 human-written TOEFL essays (year≤2020) from a Chinese educational forum (<https://toefl.zhan.com/>). The TOEFL (Test of English as a Foreign Language) is a standardized test that measures the English language proficiency of non-native speakers.

US College Admission Essays

We assembled a total of 70 authentic essays for our analysis, with 60 essays sourced from <https://blog.prepscholar.com/college-essay-examples-that-worked-expert-analysis> and 10 essays from <https://www.collegeessayguy.com/blog/college-essay-examples>.

Scientific Abstracts

We gathered a total of 145 authentic course project titles and abstracts from Stanford’s CS224n: Natural Language Processing with Deep Learning, Winter 2021 quarter (<https://web.stanford.edu/class/archive/cs/cs224n/cs224n.1214/project.html>). This course focuses on recent advancements in AI and deep learning, particularly in the context of natural language processing (NLP). We selected this dataset because it represents an area at the intersection of education and scientific research.

Statistical test

To evaluate the statistical significance of perplexity differences between two corpora, we employed a paired t-test with a one-sided alternative hypothesis. This analysis was conducted using the Python SciPy package. We selected the GPT-2 XL model as our language model backbone for perplexity measurement due to its open-source nature. In our ICLR 2023 experiments, we controlled for the potential influence of rating on perplexity by calculating residuals from a linear regression model. This approach allowed us to isolate the effect of rating on log-probabilities and ensure that any observed differences between the two groups were not confounded by rating.