

Detecting AI-Generated Writing Using GPTZero

Karen Pullet
pullet@rmu.edu

Jamie Pinchot
pinchot@rmu.edu

Evan Kinney
kinney@rmu.edu

Tyler Stewart
stewartty@rmu.edu

Computer Information Systems Department
Robert Morris University
Moon Township, PA 15108 USA

Abstract

Generative AI tools such as ChatGPT are now in widespread use and are often utilized by students to help in creating writing assignments intended to be written entirely by the student. This has spurred the need for AI detection tools such as GPTZero. This study sought to determine the accuracy of GPTZero's AI detection in identifying whether writing was created by a human, generated by AI, or a mix of both. Because many students now submit work that is a mix of both their original writing and AI-generated text, it has become more important to be able to accurately identify mixed-generated writing. The study analyzed 500 writing samples of human, AI, and mixed origin and utilized GPTZero's Deep Analysis to identify writing origin sentence-by-sentence in the mixed samples. Results from this study indicated that GPTZero accurately identified the writing origin of all samples, within an 89% to 93% accuracy rate of mixed-generated writing, and a 95-99% accuracy of writing that was written by a human or entirely by AI.

Keywords: artificial intelligence, GPTZero, ChatGPT, plagiarism detection, large language models, generative AI, plagiarism

1. INTRODUCTION

Generative AI large language models such as ChatGPT (OpenAI, 2023) can be used to create high quality human-like written work including papers, letters, reports, essays, resumes, song lyrics, poems, slide decks, answers to questions, and even computer code. The applications of ChatGPT are endless (Sadasivan, 2024) and

continue to be explored. The rate by which these tools are being adopted, especially by students, is quite rapid (Johnston et al., 2024; Kyaw, 2023; Nam, 2023).

The widespread use of these AI tools can create many conveniences and productivity benefits for writers, but it also creates risks of misuse in terms of data security, intellectual property, and ethics. The potential for misuse includes generating fake

news, fake product reviews, or otherwise manipulating web content for social engineering (Sadasivan, 2024). Ethical concerns are especially relevant in the academic environment where unauthorized use of AI for a student assignment constitutes academic cheating or plagiarism (Adeshola, 2023; Zhang et al., 2024; Sadasivan, 2024).

In some cases, the use of generative AI tools for assignments in an educational setting may be allowed, or even encouraged, by instructors in order to incorporate these new tools into students' learning experiences (Tossell et al., 2024). After all, once students enter the workforce, they will be expected to have AI literacy (Ng, 2021).

However, there are still situations when AI-tools will not be allowed for writing assignments; For instance, when the student is required to write original research, or original creative content such as a short story or poem, or is learning a writing technique or grammar. In these and other situations, the learning goals of an assignment will not be met if a student is assisted by AI. So, it has now become an issue for educators to be able to identify when student writing is original, and when it has been generated by AI.

Studies have shown that there are many difficulties for human evaluators in distinguishing between text that has been written by a human and text that has been generated by AI (Casal & Kessler, 2023; Liu et al., 2024). Traditional plagiarism detectors have also been found to be unreliable in detecting AI-generated text (Lo, 2023).

Thus, there was a need for a new class of AI detection tools that can detect whether text has been AI-generated (Zhang et al., 2024). However, many AI detection tools have been reported to be inaccurate, noting a high number of false positives (Steponenaite & Barakat, 2023; D'Agostino, 2023). Notably, OpenAI, the creator of ChatGPT, stated that AI detectors "have not been reliable enough given that educators could be making judgments about students with potentially lasting consequences" (OpenAI, 2024).

Because AI detection tools are in their infancy and little is known about whether the results provided by these tools are accurate and can be trusted, further study of them is required. One of the most popular AI detection tools is GPTZero (GPTZero.me). This tool was selected for study by the researchers to further analyze its'

accuracy. The study explores the following research question:

RQ1: How accurate are GPTZero's detection methods in determining if text was written by a human, artificial intelligence or a combination of both?

2. LITERATURE REVIEW

The use of AI in education provides new opportunities and challenges for both students and educators. It has caused mixed feelings among educators, as some view it as an opportunity for advanced teaching and learning while others focus on the drawbacks which include plagiarism, generation of incorrect information, and biases in data training (Baidoo-Anu & Ansah, 2023; Rasul et al., 2023).

Rasul et al. (2023) identified five benefits of the use of AI in higher education, which include: personalized feedback for students, facilitation of adaptive learning, support for research and data analysis tasks, development of more innovative assessments, and automation of administrative services. The authors also identify five challenges, including: academic integrity concerns, issues with reinforcement of skills sets for graduates, reliability issues, problems with assessing learning outcomes effectively, and potential biases. They argue that educators must approach the use of AI tools such as ChatGPT for academic purposes with caution in order to ensure responsible and ethical use.

Plagiarism in academia continues to be an issue that has only gained new complexities as generative AI has grown in popularity. There are challenges associated with human evaluation of writing to determine whether it is AI-generated or written by a human (Cassal & Kessler, 2023; Liu et al., 2024).

Cassal and Kessler (2023) recruited 72 reviewers and 27 editors for academic journals and had each participant complete a judgment task to determine whether paper abstracts were human or AI-generated. Though the reviewers employed a variety of techniques to judge the writing samples, they only correctly identified 38.9% of the writing that was generated by AI. Further, the study identified varying beliefs by the editors in regard to whether or not there are ethical uses of AI tools for academic research (Cassal & Kessler, 2023). Liu et al. (2024) studied 155 faculty members, researchers, and graduate students to see if they could correctly identify

writing samples that were generated by AI. The overall accuracy rate for detection by the participants in the study was 48.82%. This accuracy rate is slightly worse than random guesses (Liu et al., 2024).

In response to this issue, a number of software tools have been developed to counteract and detect AI-generated work. One of these tools is the OpenAI classifier, created by the developers of ChatGPT, which achieved 26% accuracy in a test conducted by Elkhatat et al. (2023). Another tested tool was Copyleaks, which can integrate with several key Learning Management Systems and APIs. Interestingly, this tool yielded a 99% accuracy when identifying AI-generated writing. To evaluate these tools, Elkhatat et al. (2023) utilized ChatGPT to create two 15-paragraph writing samples. To diversify the sample pool, 5 sets of human-created samples were factored in. The result ended up being that each tool has a variable difference in accuracy for identifying AI-generated writing samples. The most notable facet of this result is that the complexity of delineating human and AI-generated writing is becoming increasingly difficult. This can clearly be observed as some of the text searching tools to check for AI-generation in writing cannot detect it (Elkhatat et al., 2023).

In another analysis of AI-generated writing detectors, 16 publicly available tools were tested to observe their accuracy rates. An initial review of the 16 tools found that GPT - 2/RoBERTa, TurnItIn, and ZeroGPT were the most consistently accurate. It is important to note, however, that there was a lack of consistency across the analyses. Seventeen analyses ran undergraduate writing or short responses while the remaining 10 analyses utilized a variety of text such as abstracts, admissions essays, and exam essays (Walters, 2023). In order to successfully test the 16 tools, a sample set of documents was created. Of this sample set, 42 short papers were generated via ChatGPT. Another set of 42 documents were pulled from the Manhattan College English course 110 to add a low-level human-generated text sample (Walters, 2023). To evaluate the writing samples, the outputs of the detectors were collected and classified based on key characteristics such as whether there were numeric values in the text, whether the wording was casual or formal, and what degree of confidence was present in the writing. The results indicated that 2 of the 16 detectors correctly identified the human or AI-generated status of the documents with no incorrect or uncertain responses: Copyleaks and TurnItIn. Among the remaining detectors, it is

notable that the accuracy was fairly high, ranging around 63-88% (Walters, 2023).

Liu et al. (2024) developed a deep learning-based detection tool called CheckGPT to determine whether any given text snippet was generated by ChatGPT. They conducted a benchmarking text with 2.385 million samples of human-written and ChatGPT-written writing. Some of the samples were also not fully written by ChatGPT, but were completed or edited by ChatGPT (from writing started by a human). Results showed that CheckGPT is "highly accurate, flexible, and transferable" (Liu et al., 2024, p. 13).

Sadasivan et al. (2024) found that while some AI detectors can be accurate for basic detection, there are methods that can be used to fool the detectors into a false positive result. They develop a recursive paraphrasing attack that can be applied to text written by AI that can break a number of AI detectors. In particular, this technique can be effective in hiding AI-generated text from retrieval-based detectors and also tools that use watermarking schemes, neural networks, and zero-shot classifiers.

However, one of the most popular AI detectors, GPTZero, claims to already be implementing detection for specific paraphrasing models and maintains an updated 'greylist' of bypasser methods, which they claim to patch within days of identification (Tian, 2023). Based on its' popularity as well as these claims in regard to accuracy, the researchers chose the GPTZero AI detection tool to explore for this study.

How Does GPTZero Work?

In January of 2023, online software called GPTZero, developed by Edward Tian and Alex Cui, was launched as a response to concerns about AI-generated material (Tian, et al., 2024). GPTZero is an AI detector which checks to see if a document was created using a large language model such as ChatGPT. GPTZero detects AI on sentence structure, paragraph, and document level. GPTZero detects if tools such as ChatGPT, GPT3, Google-Gemini, LLAMA, or newer AI models were used to create a document or if it was written by a human (Tian, et al., 2024). The accuracy of GPTZero continues to increase as more texts are submitted to the model. "By learning from existing generative AI models, the tools calculate and predict the probability of words in an AI-generated sentence" (Shrivastava, para 4, 2023). GPTZero analyzes patterns of writing using syntax and sentence length to identify text created by machine learning.

GPTZero uses seven machine learning components to determine the probability of the use of AI in text: Education Module, Burstiness, Perplexity, GPTZeroX, GPTZero Shield, Internet Text Search and Deep Learning (Tian, et al., 2024). Each component provides a weighted score to the Document Classification and Document Breakdown that calculates an estimation of the amount of human-generated, AI-generated or mixed-generated writing that has been used. The Education component runs the input text against other human-written text created by students. The Burstiness component analyzes the text to see if there are patterns in the writing, whereas, Perplexity determines which words might come after one another. GPTZeroX is a component that is able to provide sentence-by-sentence classifications for human-generated, AI-generated, and mixed-generated text. GPTZero Shield defends against other tools looking to exploit the AI detector. The Internet Text Search component analyzes direct quotes from existing websites through May 2023 at the time of this writing. Lastly, the Deep Learning component is used to detect the usage of AI. Human-generated text is continuously fed into GPTZero so that it is constantly learning patterns to help determine the likelihood of AI produced material. The model is trained from creative writing, scientific writing, blogs, news articles, and more. The submissions are tested against a large-scale dataset of human- and AI-generated material (Tian, et al., 2024).

GPTZero provides three options to analyze documents. Users can copy and paste text from a source directly into GPTZero. Second, users can upload documents that are in Microsoft Word or other applications directly into GPTZero. Lastly, users can use Origin Google Chrome and Microsoft Word extensions which provide direct analysis while creating a document or reviewing a website (Tian, et al., 2024).

GPTZero has both paid and unpaid subscription plans. The tool works the same regardless of the plan used. The difference between the two versions is the number of characters that can be scanned by GPTZero. The free version only allows for seven submissions per day, limited to 5000 characters for each submission. GPTZero tracks the IP address to determine if the limit has been reached regardless of which browser is being used. In addition to the free version there are three paid subscription plans. The Essential Plan allows users up to 150,000 words per month which is equivalent to 300 pages of scanning. The Premium Plan allows 300,000 words per month which would include 600 pages of scanning to

include the AI deep scanner and is multilingual. Lastly, the Professional Plan allows for 500,000 words equivalent to 1000 pages of text, everything listed above, and allows for 10,000,000 words coverage and military grade data security (Tian, et al., 2024). All three versions use the same model to detect if text was written by a human, AI or a combination of human and AI. Each gives a breakdown for detecting the use of AI in written material with descriptions of the results.

Deep Analysis Explained

Using Deep Analysis is a feature of GPTZero which allows for a more comprehensive breakdown of the text. When scanning text, GPTZero has an output which indicates the percent of text written by a human, the percent generated by AI and will also show a percentage for a mix of human- and AI-generated content. GPTZero provides an overall probability of AI use and highlights which sentences are likely to be AI-generated. Deep Analysis quantifies the impact that each sentence makes on the overall AI probability in the writing.

3. METHODOLOGY

GPTZero claims to have a 99% accuracy rate compared to its competitors such as Originality and ZeroGPT (Tian, et al., 2024). In order to test the 99% accuracy rate, the researchers created an experiment using their own writing samples and analyzed it using GPTZero. The human samples were created from the authors own writings prior to the creation of artificial intelligence tools such as ChatGPT in November of 2022. The artificial intelligence samples were created by using ChatGPT and Claude and the mixed samples were a combination of the authors writing and the AI created samples. Due to the size limitations of the free version of GPTZero, the researchers purchased the Professional Plan which allows 500,000 words of scanning. In order to test the validity of the researchers' samples, a pilot test was first conducted. Researcher A created a sample of 20 papers. The papers were created using prior work written by the research team.

The 20 papers that were created included 6 written by a human, 7 generated using AI, and 7 that were a mix of human- and AI-generated text. The samples were numbered 1 through 20 and were then given to Researcher B to run through GPTZero. The papers given to Researcher B were anonymous and only displayed numbers. Researcher B had no idea the origin of each paper. After running the 20 papers through

GPTZero, Researcher B compared the results from the blind test to those of the known test with Researcher A's information. The 6 papers that were human-generated and the 7 papers created by AI resulted in a 98% accuracy rate with a 2% false positive rate. The results of the 7 papers that had a mix of human- and AI-generated material had an accuracy rate of 70%. The researchers realized that the reason for the percentages being off by 30% had to do with the word count in the document. When the documents that had a mix of AI and human writing were copied into the document, there were numerous font sizes. The researchers were essentially guessing the percentage instead of taking the overall word count of the document and dividing it by the percentage of AI writing and the percentage of human writing. This lesson learned from the pilot was then corrected and implemented the current study.

In addition to the 20 papers created for the pilot, the researchers conducted a preliminary study where 100 unique samples were created correcting the issues from the pilot test (Paullet, et.al, 2024). Results from the 100 samples were used as a comparison for this study in which 500 additional samples were created to analyze using GPTZero.

The writing samples for both the preliminary 100-sample study (Paullet, et.al, 2024) and the current study of 500 samples were created using the same method as the pilot test. Researcher A created 168 human samples from work written by the authors prior to November 30, 2022, 167 samples were created using the AI tool ChatGPT or Claude, and 165 samples were created combining both human and AI writing. However, in the current study, the researchers additionally utilized Deep Analysis to analyze the text line by line, highlighting which text was written by a human and which it identified as written by AI. A deep analysis was not done on the initial 100 sample study. Instead, the researchers only looked at the percentage output of whether the sample was created by a human, AI or a mix of both (Paullet, et.al, 2024). The deep analysis used in the current study provides additional insight into understanding how GPTZero identifies mixed-generated text.

The 500 writing samples created by Researcher A were anonymized and numbered 1 through 500 and put into a Microsoft Excel spreadsheet. The samples were then given to Researcher B to run through GPTZero. GPTZero allows for a single document to be uploaded or numerous documents at one time. Researcher B uploaded

the 500 documents at one time and ran a report showing the results of the experiment. A deep analysis/scan was run on all of the samples to see if GPTZero could identify line-by-line which text was written by a human, generated using AI, or contained a mix of both human and AI content. The results were then compared with the known writing origin of the samples. The results of this study are discussed below.

4. RESULTS

This study tested the functionality and accuracy of GPTZero in identifying human-generated, AI-generated, and mixed-generated text utilizing 500 unique writing samples. The research question for this study was:

RQ1: How accurate are GPTZero's detection methods in determining if text was written by a human, artificial intelligence or a combination of both?

All (100%) of the 500 samples were correctly identified, matching the known writing origin of each sample, and indicating that GPTZero accurately identified all human-generated, AI-generated, and mixed-generated samples. The number of samples created for each type is shown in Table 1.

Sample Type	Number Created
Human	168
AI	167
Mix of Human & AI	165

Table 1: Breakdown of samples created

The initial analysis of the samples identified which samples were written by a human, AI, or a mix of both. Table 2 lists the results of the unknown samples run through GPTZero by Researcher B compared to the known samples created by Researcher A. Of the 168 human samples, 167 showed a 99% chance of being written by a human and 1 sample had an output of 95-98% chance of being written by a human. The 167 AI-generated samples showed a 99% percent chance of being created using AI in 163 of the samples and 4 samples resulted in a 95-98% chance of being AI-generated. The final set of results were for samples where a combination of both human writing and writing created using AI were combined. There was a total of 165 mixed samples. Of the 165 samples that had both AI and human content, 106 had a 99% accuracy to the samples created by Researcher A. For example,

in one of the samples created with 65% of the writing being human and 35% created by AI, GPTZero was able to determine the breakdown within 1%. The percentages in the known sample were calculated by taking the total number of words in the document divided by the number of words written by a human and then the total of the AI created content to determine the percentages of the sample. There was a total of 165 mixed samples of human and AI content in which 106 samples had a 99% accuracy rate, with 43 of the samples showing a 95-98% accuracy rate, and 16 samples more than 5%. The 16 samples that showed an accuracy rate of more than 5% were examined further. It is interesting to note that the 16 samples that showed an accuracy rate higher than 5% had sentences that were intertwined in the writing of the same paragraph. Even with the writing going back and forth from human to AI, only 7 of the samples were actually over 10% (90% accuracy rate) with the highest being a 19% percent difference from the unknown to the known. Even at 81% it yielded a high enough result showing what was human-generated compared to AI-generated.

Sample Type	Number Created	1%	2-5%	More Than 5%
Human	168	167	1	0
AI	167	163	4	0
Mix of Human & AI	165	106	43	16

Table 2: Comparison of GPTZero’s identification of writing origin to known writing origin of samples

The results of the current 500 sample study are very comparable to the results of an earlier study conducted using only 100 samples (Paullet et.al, 2024). In both studies, GPTZero was 100% able to accurately identify if a sample was created by a human, AI, or a mix in all samples created. Additionally, none of the 100 samples from a previous study yielded anything lower than 82% accuracy rate and that was only in two samples. The other 98 samples were within 95-100% accuracy comparing the known to the unknown. The current study of 500 samples yielded similar results. Figure 1 shows how the data is broken down in GPTZero after a document has been scanned for detection.

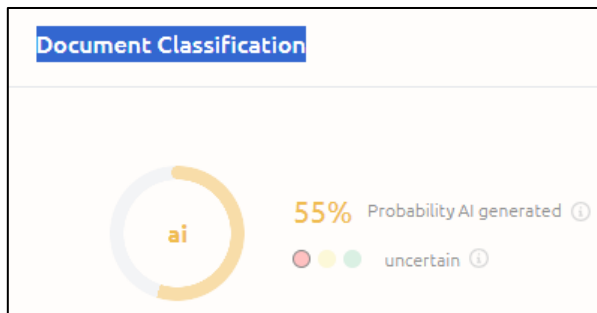


Figure 1: GPTZero Document Classification

Figure 2 is the probability breakdown of the sample showing what percentage was written by a human, AI or mixed after running it through GPTZero.

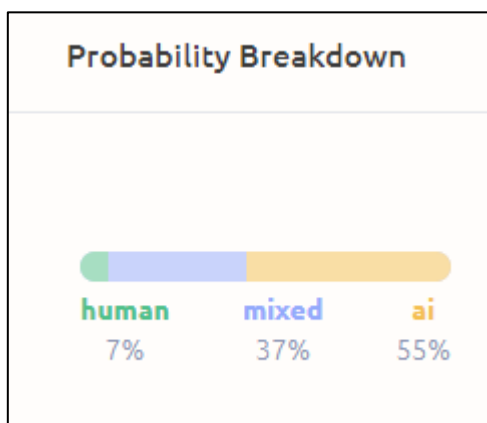


Figure 2: GPTZero Probability Breakdown

The current study of 500 samples delved deeper into the scans to see if GPTZero was able to actually detect sentence-by-sentence what was created by a human and what was created by AI. There were 165 samples that were written with a combination of human and AI that were analyzed using the Deep Scan. In all 165 of the samples, GPTZero was able to detect sentence-by-sentence what was human-generated and was created by AI. In 142 samples, GPTZero was 89% accurate while the remaining 23 samples had a 93% sentence-by-sentence accuracy rate. The Deep Scan highlights the top sentences driving the probability of AI and the top sentences driving the probability of being written by a human. Figure 2 shows a probability breakdown of being 7% created by a human, 55% AI generated and 37% is a mix of human and AI generated content. GPTZero highlights the document in different colors and provides a score along with an explanation on why it detects each probability sentence-by-sentence. "The per-sentence breakdown quantifies how much each sentence

contributes to the model's overall AI probability score. Higher scores mean greater impact on the model's prediction" (Tian, et al., 2024). The deep scan explains the highlighting colors used by GPTZero. For instance, green highlighting implies more human-like content whereas orange highlighting implies more AI-like content. The user can then hover over each highlighted line to get a more detailed prediction of the content.

5. DISCUSSION

The results of the study answer the research question indicating that GPTZero accurately identified the writing origin of all samples, within an 89% to 93% accuracy rate of mixed-generated writing, and a 95-99% accuracy of writing that was written by a human or entirely by AI. GPTZero is a reliable tool in detecting AI generated writing and can be an important tool for educators to use in classes where students are required to submit original writing without the use of AI.

There are tools claiming to be able to by-pass AI detectors. The creators of GPTZero have tested one of the AI by-pass detection tools. The result of their study that paraphrasing detection is very possible and GPTZero has already implemented tools to detect if paraphrasing of AI generated text can still be detected into the software (Tian, 2023, para 2).

Artificial intelligence is not going away. As educators, it is important to embrace the technology and find strategies that can help to eliminate misuse of AI in the classroom. Students must be taught about the risks of using AI and learn how to express themselves in their writing. They must still be able to communicate when AI is not available. Below are tips that can be utilized in the classroom to create assignments that will make it more difficult for students to use AI tools improperly and some that will aid in helping students to master AI tools when they are appropriate to use for an assignment:

1. Create assessments that cannot be answered by AI such as:
 - a. Have students write about personal experiences or discuss their learning experience in the class
 - b. Ask students to critique the default answers created by tools such as ChatGPT to questions that were created for the assignment
 - c. Require students to cite primary sources to back up their opinion or claims

- d. Require students to write about current events, which are not (at least currently) known to tools such as ChatGPT
2. Require students to produce more than one draft of their original work
 3. Ask students to create videos, podcasts, or slideshows with audio recordings
 4. Let students know that their work will be checked by an AI detector (Tian, et al. 2024).

6. LIMITATIONS

This research sought to test whether GPTZero could accurately detect human-generated, AI-generated, or mixed-generated writing samples. A limitation to the study is that only GPTZero was used to analyze the samples. Future studies should compare results from GPTZero to other similar products such as ZeroGPT or Detect GPT.

7. CONCLUSIONS

This study sought to determine if GPTZero can accurately detect if writing samples were created by a human, AI or a combination of both. Based on the outcomes of this experiment, the researchers determined that GPTZero, at the time of this research, July 2024, can accurately be used to make this determination. It is important to note that this could change at any time dependent on the advancements of AI and how quickly the developers can update GPTZero to keep up with the newest AI developments.

8. REFERENCES

- Adeshola, I. (2023). The opportunities and challenges of ChatGPT in education. *Interactive Learning Environments*, 1–14. <https://doi.org/10.1080/10494820.2023.2253858>
- Baidoo-Anu, D., & Ansah, L.O. (2023). Education in the era of generative artificial intelligence (AI): Understanding the potential benefits of ChatGPT in promoting teaching and learning. *Journal of AI*, 7(1), 53-62. <http://dx.doi.org/10.2139/ssrn.4337484>
- Casal, J.E., Kessler, M. (2023, December). Can linguists distinguish between ChatGPT/AI and human writing? A study of research ethics and academic publishing. *Research Methods in Applied Linguistics*, 2(3), 100068. <https://doi.org/10.1016/j.rmal.2023.100068>
- D'Agostino, S. (2023, June). Turnitin's AI detector: higher-than-expected false

- positives. Inside Higher Ed. <https://www.insidehighered.com/news/quick-takes/2023/06/01/turnitins-ai-detector-higher-expected-false-positives>
- Elkhatat, A.M., Elsaïd, K., & Almeer, S. (2023). Evaluating the efficacy of AI content detection tools in differentiating between human and AI-generated text. *International Journal for Educational Integrity*, 19(17), 1-16. <https://doi.org/10.1007/s40979-023-00140-5>
- Johnston, H., Wells, R., Shanks, E., Boey, T., & Parsons, B. (2024). Student perspectives on the use of generative artificial intelligence technologies in higher education. *International Journal for Educational Integrity*, 20(2), 1-21. <https://doi.org/10.1007/s40979-024-00149-4>
- Kyaw, A. (2023, December 14). *Report: Almost half of high school students use AI for schoolwork*. Diverse Education. <https://www.diverseeducation.com/student-issues/article/15660259/report-almost-half-of-high-school-students-use-ai-for-schoolwork>
- Liu, Z., Yao, Z., Li, F., Luo, B. (2024, March). On the detectability of ChatGPT content: Benchmarking, methodology, and evaluation through the lens of academic writing. <https://doi.org/10.48550/arXiv.2306.05524>
- Lo, C.K. (2023). What is the impact of ChatGPT on education? A rapid review of the literature. *Education Sciences*, 13(4), 410. <https://doi.org/10.3390/educsci13040410>
- Nam, J. (2023, November 22). *56% of college students have used AI on assignments or exams*. Best Colleges. <https://www.bestcolleges.com/research/most-college-students-have-used-ai-survey/>
- Ng, D.T.K., Leung, J.K.L., Chu, S.K.W., & Qiao, M.S. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2(2021), 100041. <https://doi.org/10.1016/j.caeai.2021.100041>
- OpenAI. (2023, March). GPT-4 technical report. <https://cdn.openai.com/papers/gpt-4.pdf>
- OpenAI. (2024, July). How can educators respond to students presenting AI-generated content as their own? <https://help.openai.com/en/articles/831335>
- 1-how-can-educators-respond-to-students-presenting-ai-generated-content-as-their-own
- Paullet, K, Pinchot, J., Kinney, E., & Stewart, T. (2024). Can GPTZero detect if students are using artificial intelligence to create assignments? IACIS 2024 Conference Proceedings.
- Rasul, T., Nair, S., Kalendra, D., Robin, M., de Oliveira Santini, F., Laderia, W.J., Sun, M., Day, I., Rather, R.A., & Heathcote, L. (2023). The role of ChatGPT in higher education: Benefits, challenges, and future research directions. *Journal of Applied Learning & Teaching*, 6(1), 1-16. <http://journals.sfu.ca/jalt/index.php/jalt/index>
- Sadasivan, V.S., Kumar, A., Balasubramanian, S., Wang, W., & Feizi, S. (2024, February). Can AI-generated text be reliably detected? <https://arxiv.org/pdf/2303.11156>
- Shrivastava, R. (2024). With seed funding secured, AI detection tool GPTZero launches new browser plugin. *Forbes*. <https://www.forbes.com/sites/rashishrivastava/2023/05/09/with-seed-funding-secured-ai-detection-tool-gptzero-launches-new-browser-plugin/?sh=1ef814b33f7b>
- Tian, E. (2023). Ways to by-pass AI detection? <https://gptzero.me/news/gptzero-by-passers>
- Tian, E., & Cui, A. (2024). GPTZero: Towards detection of AI-generated text using zero-shot and supervised methods.
- Steponenaite, A., & Barakat, B. (2023, July). Plagiarism in AI empowered world. In *International Conference on Human-Computer Interaction*, 434-442. https://sure.sunderland.ac.uk/id/eprint/15622/7/Conference_paper_template_AI_manuscript_edit2.pdf
- Tian, E. (2023, March). Ways to by-pass AI detection? GPTZero. <https://gptzero.me/news/gptzero-by-passers>
- Tossell, C.C., Tenhundfeld, N.L., Momen, A., Cooley, K., & de Visser, E.J. (2024). Student perceptions of ChatGPT use in a college essay assignment: Implications for learning, grading, and trust in artificial intelligence. *IEEE Transactions on Learning Technologies*, 17, 1069-1081. <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=10400910>

Walters, W.H. (2023). The effectiveness of software designed to detect AI-generated writing: A comparison of 16 AI text detectors. *Open Information Science*, 7(1). <https://doi.org/10.1515/opis-2022-0158>

Zhang, Y., Ma, Y., Liu, J., Liu, X., Wang, X., & Lu, W. (2024, April). Detection vs. anti-

detection: Is text generated by AI detectable? *International Conference on Information*, 209-222. https://doi.org/10.1007/978-3-031-57850-2_16