

ChatGPT in the Classroom:

Uses, Limitations, and Student and Teacher Experiences



The education world has entered the first full school year, 2023-2024, with generative Artificial Intelligence (AI). It was in early 2023, the middle of last school year, when the introduction of ChatGPT sparked a wave of apprehension among teachers. Fears were rampant that this advanced technology would disrupt the educational landscape. As we navigate through the first full year of its implementation, we begin to see its potential benefits and understand its potential drawbacks. Educators continue to explore and better understand the implications of this innovative technology and will continue to do so for some time into the future.

Despite initial fears, generative AI is showing promise in helping educators with their jobs and offering new opportunities to students. At the same time, the current limitations of the technology may cause a variety of unexpected consequences.

ChatGPT is not the only source for generative AI. There are dozens, if not hundreds, of similar Large Language Models (LLMs) performing variations of the same process. Some notable apps include Google's Bard, Microsoft's Bing Chat, Socrat.AI, and Anthropic's Claude. ChatGPT was the first widely known model, publicly released in November of 2022. In this

article, therefore, I use ChatGPT as a shorthand for all generative AI, even if some tools act somewhat differently.

LLMs Infiltrate Education

ChatGPT was the fastest public tech deployment in history. Facebook took two years to reach one million users. Instagram took two and a half months. ChatGPT took five days.¹ This quick implementation created tension in classrooms as teachers wanted to address its use but didn't know how, especially given the rapid developments and changes. Any thoughts of bringing it up were tempered by the possibility that they might be giving away the secret to the

BY LORIN KOCH

most efficient cheating method that has ever existed.

Teachers need not have worried about that—their students already knew. Students aren't ignorant about technological development. They pay attention to news, albeit usually through YouTube videos or social media. A former student of mine, who graduated from 12th grade in 2023, says students learned about ChatGPT online, through memes and TikTok posts, and then a popular TV episode² midway through the school year. Another former student, a 10th grader last year, says YouTube videos discussed ChatGPT as far back as November 2022. This student was an early adopter, signing up in December and spreading the word to other students in January and February of 2023.

Early reactions included excitement (by students) and horror (by teachers). Teachers and professors were concerned that ChatGPT would immediately make essay writing obsolete, as students could complete their papers in seconds with very little effort. Would this be the end of the essay? The end of homework altogether?

Before long, almost everyone knew about ChatGPT. Students at my Adventist school who were confident of their ability to write their own thoughts were perhaps less tempted, but younger or less-mature students saw it as a shortcut to avoid tasks seen as time-consuming, difficult, or annoying. One student was not proud to admit using it on assignments, but the ability to get out of doing work was too strong to ignore. This student used ChatGPT both to help with brainstorming and to write full papers, which then needed editing. If saving work was the goal, ChatGPT didn't come through: The time spent fixing the issues in the AI-generated papers was longer than the time it would have taken to write the paper.

The Possibilities of AI in Schools

As the technology became better known and its use more widespread, some teachers began to shift from hor-

ror to guarded optimism, coming to see ChatGPT as another tool in their educational toolbox. Nathaniel Whittemore, host of the podcast *The AI Breakdown*, says that educators and students are always among the first adaptive users of new technology.³ Technology has presented this type of possibility many times. To offer a comparison, the calculator disrupted the way math classes were taught, but after a period of resistance, teachers quickly learned to adapt their curricula to the new tool.⁴

ChatGPT simply has no way of fact-checking its information. Like an unprepared student giving a presentation, it is trying to use the text it knows to sound plausibly human, which sometimes results in very confident-sounding but very incorrect information.

Generative AI presents another sort of challenge. This technology was totally unknown to the majority of educators just a year ago. However, if teachers can be open to the opportunities provided by generative AI, it offers many beneficial uses. Box 1 on page 6 contains a list of ways generative AI can be used in the classroom that the author has gathered from various sources and teachers' experiences.

Many more ideas can be found online. Free idea lists exist, such as a list of more than 100 creative ideas in Zenodo.⁵ Also, researchers are beginning to publish books on the topic.⁶ Many of the ideas in the list in Box 1 came from Shah or Skrabut, among other sources. OpenAI also introduced a set of blog posts, help tools, and FAQs at the beginning of the 2023-2024 school year for educators dealing with generative AI.⁷

What Limitations Can Teachers and Students Expect?

While generative AI seems to provide impressive text most of the time, it does have a variety of shortcomings that all users should understand. Teachers need to know that ChatGPT and other AI models will often lack common sense, be inaccurate, and make up details. Students should also be aware of AI's tendency to make things up and the fact that it has very little ethical guidance. Also, currently, there seems to be limited access to high-quality peer-reviewed materials that exist behind paywalls of reputable journals.⁸

ChatGPT's issues with common sense and accuracy are both a result of the fact that the model doesn't actually understand words or language.⁹ The computer "brain" is simply predicting what it thinks you want to hear based on mathematical calculations and probability. ChatGPT breaks letters and words down into binary numbers, and it has no idea whether its information is accurate or not. This makes for some puzzling interactions, such as these:

User: *How many R's are in the word 'Idiosyncrasy'?*

ChatGPT: *There are three R's in the word 'Idiosyncrasy.'*

User: *Explain your answer. Locate all of the R's.*

ChatGPT: *I apologize for the incorrect response in my previous answer. There are actually two R's in the word 'Idiosyncrasy.' Here's the correct breakdown:*

The first R appears after the letter 'i.'

Box 1. Ways Generative AI Can be Used

Creating plans for class

- Generating lesson plans on a specific topic
- Sequencing material
- Formulating topics of study for different focus groups
- Creating interactive activities [such as science labs] on a specific concept
- Defining collaborative tasks
- Suggesting performance tasks that align with learning objectives

Generating text for classroom use

- Creating class discussion topics
- Assembling quiz questions
- Writing distractors for multiple-choice questions
- Compiling case studies for a topic being studied
- Writing samples at different reading levels
- Creating computer code for specified tasks
- Assembling PBL or STEAM projects related to a topic
- Taking existing questions and remaking them into a more interactive project
- Generating a “Jeopardy” question board for review

Differentiation

- Offering enrichment for gifted students
- Identifying tips for helping ADHD students
- Zone of Proximal Development – Sets of increasingly difficult tasks
- Generating tasks at different levels of Bloom’s Taxonomy

Ideas for non-academic areas

- Bulletin board themes
- Icebreaker activities
- Group activities
- Spiritual programming
- Generating a classroom newsletter – Tell ChatGPT to ask you a series of

questions about what’s happening in your class, and then have it write the news articles

Help with grading

- Giving immediate feedback on student understanding
- Summarizing text
- Creating rubrics
- Outlining writing
- Explaining errors in writing

Creative classroom applications

- “Respond as a literary character” – Students get into a book by having a “conversation” with one of its characters
- “Be a debating partner” – Students test their arguments by having the AI assess and poke holes in their arguments
- “When am I ever going to use this?” AI can answer this age-old student question
- Have ChatGPT write the first draft of a paper, and then give students the task of fact-checking the AI
- “Give me an array of potential perspectives” on a topic
- Use ChatGPT as a socratic dialogue tool

Help with difficult situations

- Writing reference letters for students
- Creating letters to parents that deal with a variety of topics

Reminder: Never disclose personal information to an AI chatbot. The conversations may be kept for training purposes and may be viewed by many people!

The second R appears after the letter ‘o.’

User: What is 6257×962 ?

ChatGPT: The product of 6257 multiplied by 962 is 6,022,834.

User: No it isn’t!

ChatGPT: I apologize for the error in my previous response. Let’s calculate the correct product of 6257 multiplied by 962:

$$6257 \times 962 = 6,014,134$$

So, the correct result is 6,014,134. Thank you for pointing out the mistake.¹⁰

(By the way, the correct answer is 6,019,234.)

Why would it make such a basic mistake like the number of R’s in a word, or the answer to a multiplication problem that any calculator gets correct? Remember that LLMs are trained on online text. It is very unlikely that anyone online has ever asked how many R’s are in the word *idiosyncrasy*. They wouldn’t need to: Anyone who can recognize letters can count the R’s. Similarly, while it is very likely that the multiplication problem 7×5 is on numerous websites, the problem above may not be on any Internet source. The AI has to do its best to come up with what it “thinks” you want. It’s pretty close both times! But neither answer is accurate. In the above examples, ChatGPT has very little to draw on from its training data.

This also helps explain why ChatGPT has a tendency to make up information. Observers sometimes refer to this as hallucination, but some researchers prefer the term *confabulation* to keep from over-humanizing the computer models.¹¹ ChatGPT simply has no way of fact-checking its information. Like an unprepared student giving a presentation, it is trying to use the text it knows to sound plausibly human, which sometimes results in very confident-sounding but very incorrect information. A prominent example of this was the story of the lawyer who used ChatGPT to help prepare for a case, and ended up citing six cases by name, none of which actually existed!¹² (This true story

should be mentioned to all students who might be tempted to use generative AI.)

One other point is related to the fact that text generated by LLMs is sometimes hurtful and negative. Because ChatGPT is the sum of the material that makes up its source data, it may give information and advice that could be biased or hurtful. While programmers have given LLMs the goal of being helpful and appropriate, think about all of the text you've seen on the internet. How much biased or hurtful text is online? That should give you an idea about some of what an LLM has to sift through as it tries to be helpful to humans.

AI and Christian Schools

Faith-based schools have some specific concerns and opportunities related to the use of generative AI. A major area of interest for believers is the fact that AI was created by humans. God created humans in His image and bestowed us with a spark of divine creativity, which we then used to build technology in our image. AI technology is not human consciousness, of course. But programmers want AI to respond like a human would, and this brings up significant questions about the ethics of how we want it to relate to us.

A 2022 dissertation discusses in depth the topic of ethics in AI.¹³ In this study, the author interviewed several people who work in fields that bridge the intersection of faith and technology. These interviews discuss how the ethics desired and generated by programmers, such as efficiency and accuracy, differ from Christian ethics, such as humility, altruism, sacrifice, mercy, and love. Several of the interviewees discussed the possible benefits of including love in ethical systems, compared to attempts to program ethics into AI.

Concerns mentioned earlier, such as bias and harmful content, relate directly to generative AI's ethical programming. Other major ethical considerations brought up by critics of AI include the use of deepfakes, the ability to protect private data, and

fears of turning over human decision-making to AI.¹⁴ Huizinga's interviewees suggest that including God in the ethical programming of AI could alleviate some of those concerns. When the reason for ethics is based on God first loving us, we act ethically as a way of uplifting others toward God. But when an ethical system is based on efficiency, accuracy, and *avoiding* specific content, it is likely that it will not act in ways that truly benefit people. As much as we

Imagine a classroom in which the teacher is taking an unorthodox approach to learning. Instead of posting learning objectives and key learnings, the teacher tells the students, "Do whatever comes to mind. I have a goal for you, and I'll let you know if you're getting close. Try stuff!"

may think of AI as a human intellect, it is helpful to remember that it is a tool programmed in ways that do not necessarily include love for all humankind.

As a tool, ChatGPT can still provide a lot of benefits, even specifically to the unique needs of people at Christian schools. Generative AI offers exceptionally good programming suggestions for spiritual events and services. Consider a busy school chaplain who has a Week of Prayer to plan. ChatGPT could quickly provide

cohesive themes, engaging activities, and reflection questions to connect students' coursework to concepts presented in the meetings. But there is always the concern of accuracy. While ChatGPT knows a lot about the Bible, it may not provide material that is theologically sound or accurate to the original text.¹⁵ This rule of thumb for students also applies here: Use AI for ideas but do your own work. God calls us to speak the truth in love (Ephesians 4:15). Teachers and spiritual event planners should make sure that the AI suggestions they use are in line with their beliefs, keeping in mind the ethical concerns discussed above.

What Is ChatGPT Doing?

When ChatGPT comes up with something odd, and we don't know why, it's an example of what is known as the "alignment problem."¹⁶ This is when the goals of the AI don't match up with our goals for it. In these situations, there is literally no way for us humans to understand why it did what it did because it has to go through multiple trillions of calculations to provide answers. A good step in understanding why LLMs confabulate and introduce errors is to gain some understanding about what the computers are actually doing behind the scenes.¹⁷

Imagine a classroom in which the teacher is taking an unorthodox approach to learning. Instead of posting learning objectives and key learnings, the teacher tells the students, "Do whatever comes to mind. I have a goal for you, and I'll let you know if you're getting close. Try stuff!"

As the students "try stuff," the teacher keeps giving them a thumbs-down. Writing on the board gets a thumbs-down. Drawing a picture gets a thumbs-down. Opening a science textbook gets a thumbs-down.

Out of exasperation, one student gives up and heads over to the Lego corner, and the teacher, surprisingly, gives a thumbs-up! Reinvigorated, the students start grabbing Legos and putting them in different configurations. Eventually, they create a model of a

subtraction problem: 8 Legos - 5 Legos = 3 Legos. Without ever being told the objective, the students arrive at a correct answer.

This is, on a very small scale, how ChatGPT was trained. Instead of programming in rules and logical tests, the computer was asked questions. If its answers were good, humans essentially gave it a thumbs-up. In the classroom scenario above, with the time constraints of a school day, it's very possible that the students might never determine the teacher's goal. But repeat this process millions, or billions, of times, and you get a computer that is fairly good at answering in a way humans will understand. This is probably not the best way to teach a class of students. The teacher knows the students arrived at a correct answer, but doesn't really know what they understand, if anything. That is basically what happens with the answers given by ChatGPT.

Tasks LLMs Do Poorly

Looking at this whole picture, teachers may feel a sense of hopelessness about whether they can avoid ChatGPT taking over their students' homework. Some might be tempted to give up trying to teach accuracy, writing, basic comprehension, or even critical thinking since LLMs can do it so well and don't seem to "care" when they fail. This would be an understandable reaction, but one that needs to be looked at honestly.

The two teachers I interviewed for this article both said they don't think ChatGPT should take the place of traditional skills of writing, organizing, and putting concepts together. Even if generative AI can do basic tasks well, it is still important for students to be able to do them. One of the teachers mentioned that students need rudimentary skills in order to pursue higher-order skills. Having AI do those simpler tasks could cause teachers to rush to more complex skills, leaving behind students who need help with

lower-order skills and creating gaps in understanding.¹⁸

Reflecting on this idea, one teacher commented that we need to avoid all-or-nothing thinking about technology. Not all academic skills have to be done using the Internet, or even typed. Critical-thinking skills are often best learned in conversation or through activities in class. Furthermore, even if ChatGPT can write a sonnet, that doesn't mean people should stop writing sonnets.

Teachers can also structure class time to have students complete tasks in school rather than on their own, where the opportunity to use an LLM is greater.

That being said, it is still valuable to have an understanding of tasks that ChatGPT either currently can't do or struggles to do well. This list is most useful for teachers who are attempting to minimize the likelihood that students will try to use generative AI to do their classwork.

LLMs are bad at...

- Analysis of specific data from images or video shown in class: LLMs have begun to be able to understand what is in images, but still struggle to comprehend the significance of material in images and videos.
- Analysis that draws on, or cites, class discussion. Tasks that reference what happened in class can only be completed by people who were there.
- Personal reactions to specific cited sources: Teachers can ask students to respond to, and evaluate, specific material with their own understanding. LLMs can "pretend" to

apply information to their "lives," but a real-life connection will be much more vivid.

- Themes: Because they don't have a sense of overall cohesion, LLMs struggle with big-picture analysis, and with using a smaller illustration of a larger concept. They tend to repeat themselves, or even contradict themselves, on longer passages.

In addition to the points above, teachers can also structure class time to have students complete tasks in school rather than on their own, where the opportunity to use an LLM is greater. For instructors looking to teach important writing skills, pointers for use of class time include:

- Assign in-class writing;
- Have students brainstorm and pre-write by hand. (Requiring students to hand-write essays and longer passages is not recommended.¹⁹ Additional adaptations might need to be made for students with other learning needs);
- Break down writing into smaller steps that can be completed in class;
- Require the use of some specific sources to which every student has access;
- Hold individual writing conferences where you ask students to explain their reasoning.

Catching ChatGPT in the Act

The two Adventist teachers I talked to for this article both caught their students using ChatGPT last school year. They reflected on how strange it was to discover this. Both of them knew on a first read-through that the papers were not written by students. The language didn't reflect how the topic was discussed in class, and the work didn't "match" what the students produced earlier in the school year. Additionally, the content was described as being "thorough, past the point that most students would stop talking about something," but also vague at the same time.

When confronted with the evidence, most of the students admitted

Box 2. Empowering Education With AI: Text and Image Generation Tools in the Classroom

Adventist educators desiring to learn more about using AI in the classroom can now access the self-paced online course “Empowering Education with AI: Text and Image Generation Tools in the Classroom” through the Adventist Learning Community at <https://www.adventistlearningcommunity.com/courses/gpt-ai>.

using ChatGPT right away. One of the students claimed to only have used it to edit his writing and check grammar, although the teacher felt dubious about this assertion.

Both of the teachers I talked to reported being hesitant to bring up generative AI with their students. When teachers haven’t used ChatGPT much themselves, they may not be confident in attempting to inform others about it. One of the teachers didn’t know anything about ChatGPT before catching it in her classroom. She knew the paper wasn’t written by a student but didn’t know what it was. Putting it through a plagiarism detector didn’t help, and she felt confused.

Can You Use ChatGPT to Catch ChatGPT?

Toward the end of the 2022-2023 school year, dozens of companies began to come out with tools designed to catch AI-generated text in assignments. These tools, called “classifiers,” seemed to promise to be the new version of plagiarism checkers. OpenAI released its own classifier in January of 2023, with the warning that it was wrong a lot of the time.²⁰ It could accurately classify AI text only 26 percent of the time, which of course meant that it missed 74 percent, and wrongly classified human-written text as AI nine percent of the time!

The theory was that classifiers would get better with time, and eventually be able to tell with a high degree of certainty whether a human

or a computer wrote text. This did not prove to be true. AI-generated text is much more difficult to catch than plagiarized text, which was likely published somewhere else. The challenge is that AI-generated text doesn’t already exist anywhere, so classifiers have to look for other characteristics of text. For some comparisons, computers are more predictable than humans, usually completing sentences in the way others do, and more consistent with sentence length and structure.

With those characteristics, it might seem possible to determine whether text is human-generated. But thus far, they are not reliable indicators.²¹ A lot of computer-generated text is still unflagged (false negatives), and human-written text is often flagged (false positives). Wrong classifications have consequences. Wrongly failing students over false positives is unfair, and succeeding by passing off false negative AI-generated text as one’s own is unethical.

False positives have caused significant upheaval already, such as when a college professor threatened to fail his entire class over ChatGPT use.²² Even worse is the fact that English-language learners are more likely to have their writing falsely flagged as AI-generated, likely due to predictable writing limited by vocabulary.²³

The classification problem is so bad that some teachers might be tempted to just ask ChatGPT if it wrote something. Responding to that idea, OpenAI stated in their Teacher Guide FAQ that asking ChatGPT if it wrote something will fail: “ChatGPT has no ‘knowledge’ of what content could be AI-generated or what it gen-

erated. It will sometimes make up responses to questions like ‘did you write this [essay]?’ or ‘could this have been written by AI?’ These responses are random and have no basis in fact.”²⁴

OpenAI offers some suggestions, though, for teachers hoping to avoid the problem of students turning in AI-generated text as their own.²⁵ Suggestions include incorporating generative AI in class, teaching the students how to use it ethically and responsibly, and having students show their work throughout the process.

Conclusion

Generative AI has the potential to revolutionize the way we teach and learn. With the explosion of LLMs such as ChatGPT, we are seeing a new world of personalized learning that can help students and teachers alike. However, it’s important to remember that these models are still relatively new and unpredictable. Their limitations need to be understood as teachers (and students) use them in classes.

As we move forward integrating AI into our classrooms, we would do well to remember the ethical considerations involved. Teachers and students alike need to be aware of the limitations of AI, including confabulation, inaccurate information, and the potential for bias and questionable content. Teachers need to take steps to mitigate these risks and educate their students.

Ultimately, the success of generative AI in education will depend on how well we can balance the benefits with the limitations. By working together to develop best practices and guidelines, and sharing our successes, we can ensure that this technology is used responsibly and effectively. ✍️

This article has been peer reviewed.



Lorin Koch, EdD, is a course author for Griggs International University, a course developer and instructor for Adventist Learning Community (ALC), and a curriculum development associate for Pedagogy.Cloud, an educational consulting firm. With 19 years of experience in Adventist education, Dr. Koch has worked as a secondary school teacher and registrar at Livingstone Adventist Academy (Salem, Oregon, U.S.A.) and Indiana Academy (Cicero, Indiana, U.S.A.). He has an MAT in Curriculum and Instruction from Walla Walla University (College Place, Washington, U.S.A.), and an EdD in Curriculum Studies from the University of South Carolina (Columbia, South Carolina, U.S.A.). This article emerged out of his presentation at the North American Division Educators Convention in Phoenix, Arizona, in August 2023, and out of his course development work for ALC.

Recommended citation:

Lorin Koch, “ChatGPT in the Classroom: Uses, Limitations, and Student and Teacher Experiences,” *The Journal of Adventist Education* 85:3 (2023): 4-10. <https://doi.org/10.55668/jae0046>.

NOTES AND REFERENCES

1. Noam Hassenfeld (Host), “The Black Box: In AI We Trust?” [Audio podcast episode]. In *Unexplainable*, Vox Media Network (July 19, 2023): <https://pod.link/unexplainable/episode/4517829cfa939b38f5464d865325c24f>.
2. Jez Corden, “The Latest Episode of South Park is About . . . ChatGPT,” *Windows Central* (March 7, 2023): <https://www.windowscentral.com/microsoft/the-latest-episode-of-south-park-is-about-chatgpt>.
3. Nathaniel Whitemore (Host), “Are Educators Ready for a ChatGPT School Year?” [Audio podcast episode]. In *The AI Breakdown*. The Breakdown Network (September 1, 2023): <https://pod.link/1680633614/episode/dd5a5e33112b0522483800a238c9f00e>.

4. William Pang, “The Common High-School Tool That’s Banned in College,” *The Atlantic* (December 22, 2016): <https://www.theatlantic.com/education/archives/2016/12/the-conundrum-of-calculators-in-the-classroom/493961/>.
5. Chrissi Nerantzi et al., eds., “101 Creative Ideas to Use AI in Education: A Crowdsourced Collection,” *Zenodo* (2023): <https://zenodo.org/record/8224168>.
6. Priten Shah, *AI and the Future of Education: Teaching in the Age of Artificial Intelligence* (San Francisco: Jossey-Bass, 2023); Stan Skrabut, *80 Ways to Use ChatGPT in the Classroom* (Self-Published, 2023).
7. OpenAI, “Teaching With AI” (August 31, 2023): <https://openai.com/blog/teaching-with-ai>.
8. Sara Guaglione, “Why Protecting Paywalled Content From AI Bots is Difficult Business” (August 7, 2023): <https://digiday.com/media/why-protecting-paywalled-content-from-ai-bots-is-difficult-business/>; Emily Dreibelbis, “Browse With Bing’ Disabled on ChatGPT Plus Because It Bypassed Paywalls” (July 5, 2023): <https://www.pcmag.com/news/browse-with-bing-disabled-on-chatgpt-plus-because-it-bypassed-paywalls>.
9. Paul Thagard, “Why Is ChatGPT So Smart and So Stupid?” *Psychology Today* (February 23, 2023): <https://www.psychologytoday.com/us/blog/hot-thought/202302/why-is-chatgpt-so-smart-and-so-stupid>.
10. Conversations generated by the author, using ChatGPT’s model GPT 3.5.
11. Benj Edwards, “Why ChatGPT and Bing Chat Are So Good at Making Things Up: A Look Inside the Hallucinating Artificial Minds of the Famous Text Prediction Bots,” *Ars Technica* (April 6, 2023): <https://arstechnica.com/information-technology/2023/04/why-ai-chatbots-are-the-ultimate-bs-ma-chines-and-how-people-hope-to-fix-them/>.
12. Kathryn Armstrong, “ChatGPT: US Lawyer Admits Using AI for Case Research,” *BBC News* (May 27, 2023): <https://www.bbc.com/news/world-us-canada-65735769>.
13. Gretchen Huizinga, *Righteous AI: The Christian Voice in the Ethical AI Conversation* (PhD dissertation, University of Washington, 2022): https://digital.lib.washington.edu/researchworks/bitstream/handle/1773/48925/Huizinga_washington_0250E_24261.pdf?sequence=1#:~:text=This%20study%20adds%20critical%20insights,intelligence%20is%20designed%20and%20implemented.
14. Selin Akgun and Christine Greenhow, “Artificial Intelligence in Education: Addressing Ethical Challenges in K-12 Settings,” *AI Ethics* 2 (2022): 431-440. <https://doi.org/10.1007/s43681-021-00096-7>.
15. Jennifer A. Kingston, “Religious Leaders Experiment With ChatGPT Sermons,”

Axios (March 10, 2023): <https://www.axios.com/2023/03/10/pastors-chatgpt-sermons-rabbi-minister>.

16. Edd Gent, “What is the AI Alignment Problem and How Can it be Solved?” *New Scientist* (May 10, 2023): <https://www.newscientist.com/article/mg25834382-000-what-is-the-ai-alignment-problem-and-how-can-it-be-solved/>.
17. Samuel R. Bowman, “Eight Things to Know About Large Language Models” (Unpublished manuscript 2023): <https://arxiv.org/pdf/2304.00612.pdf>.
18. Shah, *AI and the Future of Education: Teaching in the Age of Artificial Intelligence*.
19. Anna Mills and Lauren M. E. Goodlad, “Adapting College Writing for the Age of Large Language Models Such as ChatGPT: Some Next Steps for Educators,” *Critical AI* (January 17, 2023, updated April 17): <https://criticalai.org/2023/01/17/critical-ai-adapting-college-writing-for-the-age-of-large-language-models-such-as-chatgpt-some-next-steps-for-educators/>.
20. Jan Hendrik Kirchner et al., “New AI Classifier for Indicating AI-written Text,” *OpenAI* (January 31, 2023): <https://openai.com/blog/new-ai-classifier-for-indicating-ai-written-text>.
21. Benj Edwards, “Why AI Detectors Think the US Constitution Was Written by AI,” *Ars Technica* (July 14, 2023): <https://arstechnica.com/information-technology/2023/07/why-ai-detectors-think-the-us-constitution-was-written-by-ai/>.
22. Pranshu Verma, “A Professor Accused His Class of Using ChatGPT, Putting Diplomas in Jeopardy,” *Washington Post* (May 19, 2023): <https://www.washingtonpost.com/technology/2023/05/18/texas-professor-threatened-fail-class-chatgpt-cheating/>.
23. Weixin Liang et al., “GPT Detectors Are Biased Against Non-native English Writers,” *Patterns* 4:7 (July 14, 2023): 100779–100779. <https://doi.org/10.1016/j.patter.2023.100779>.
24. OpenAI Help Center Educator FAQ, “Can I Ask ChatGPT If It Wrote Something?” (n.d.): Retrieved September 11, 2023, from <https://help.openai.com/en/articles/8318890-can-i-ask-chatgpt-if-it-wrote-something>.
25. OpenAI Help Center Educator FAQ, “How Can Educators Respond to Students Presenting AI-generated Content as Their Own?” (n.d.): Retrieved September 11, 2023, from <https://help.openai.com/en/articles/8313351-how-can-educators-respond-to-students-presenting-ai-generated-content-as-their-own>.