

2. GUIDELINES FOR MANAGING AND UTILISING GENERATIVE ARTIFICIAL INTELLIGENCE WRITING TOOLS

Justina Amakali, Raymond Greenlaw & Kauna Mufeti

COPYRIGHT
© The Author(s)
Published under a Creative
Commons Attribution 4.0
International Licence
(CC BY 4.0)

Abstract

Students use artificial intelligence (AI) generative writing tools to complete writing assignments, regardless of whether such tools are allowed or not. The quality of writing produced by AI tools is superior to that of many university students in terms of grammatical correctness, reduced spelling errors, increased average word length, expanded vocabulary, varied sentence structure, decreased word repetition, and increased punctuation usage. This research suggests ways in which AI generative writing tools can be used to improve student writing. It discusses how to make students better editors, so they can improve ChatGPT's output or tailor it to specific needs. It also provides insights on how to make students better fact checkers, so that they can eliminate hallucinations that AI writing tools generate. These errors of fact make the AI writing tools far less valuable than they would be otherwise. Our work highlights the need to focus writing exercises across the curriculum more on editing and fact-checking. We provide numerous insights into programs such as ChatGPT to help instructors without a technical background identify, understand, and process student writing generated by such programs in an improved fashion. We identify several tricks that students use to disguise writing produced by ChatGPT.

Keywords: ChatGPT, generative artificial intelligence, AI writing tools

Introduction

The Dartmouth Conference of 1956 is considered by many to be the beginning of the field of artificial intelligence (AI) (Moor, 2006). John McCarthy, Marvin Minsky, Allen Newell, Herbert Simon, and Alan Turing are the discipline's founding fathers (Abreu, 2023). Over the following 25 years, the foundations of AI developed with many new techniques emerging. Relative to modern machines, early computers had extremely limited processing power and memory capacity, so no one envisioned the types of AI apps that would be developed. The personal computer (PC) was only developed in the mid-1970s (Horn and Winston, 1975), and its memory capacity was measured in kilobytes (K), not megabytes (M), not gigabytes (G), and certainly not terabytes (T), petabytes (P), or exabytes (E). Although progress in AI continued steadily over the subsequent three decades, AI never caught the attention of the masses. No one anticipated that essentially all human writings, recordings, entertainment, economic matters, and social endeavours would be publicly available on the Web (founded in 1989,

CERN, 2023) and that computing resources would be available at such a low cost. Big data became gigantic and of unimaginable size. Moore's Law essentially came to fruition over a period of six decades (Moore, 1965). The outcome of these advances is that more progress has been made in AI in the last few years than in the previous 65.

The first spelling checker was developed by Les Earnest in 1961 (Gorin, 2009). By modern standards, it was primitive at best. Interactive spelling checkers were developed (Gorin, 2009), and more sophisticated and faster spelling checkers emerged in the 1980s for PCs (Peterson, 1980). In the 1990s and 2000s, spelling correction features were incorporated into Web browsers (Wikipedia, 2023). Modern spelling checkers are fast, check spellings "on the fly," contain huge vocabularies, allow users to augment the dictionary, provide a list of possible correct spellings, and do much more. AI writing tools next progressed to basic-style suggestions and grammar checking. Writer's Workbench was one of the first such tools (Writer's Workbench, 1986). Over the next few decades, grammar checkers, Grammarly (Grammarly, 2023), QuillBot (QuillBot., 2023), and Zoho Writer (Davis, 1986) to name but a few, steadily improved and expanded, checking not only for grammar issues, but looking at verb conjugations, analysing sentence structure, detecting plagiarism, checking tone, suggesting word choice, and so forth. The overall progress in AI writing tools was slow and incremental, but has been steady.

This situation changed drastically when a writing revolution exploded in November 2022 as OpenAI released ChatGPT (Chat Generative Pre-Trained Transformer). It became the fastest-growing app in history (Duarte, 2023). Since its introduction, the public's fascination with AI has mushroomed dramatically. ChatGPT takes a user's prompt and immediately responds with well-formed text that it generates. Although there are a host of other large language models such as BARD, Claude, Falcon, LLaMA, and PaLM, our focus in this paper is exclusively on ChatGPT. Various versions of ChatGPT are available, but for this research, we focus on the free version of ChatGPT 3.5 (ChatGPT 3.5, 2023). As an indication of what ChatGPT is capable of, we entered the following prompt into ChatGPT: 'Write three sentences about student use of ChatGPT.' Within seconds, ChatGPT responded with the following:

"Students frequently use ChatGPT to seek assistance with homework, research, and understanding complex concepts across various subjects. Its availability as a 24/7 virtual tutor and quick responses makes it a valuable resource for students looking to enhance their learning experience. Additionally, ChatGPT can also aid in brainstorming ideas and refining essays or academic papers, further supporting students in their educational pursuits."

On reading ChatGPT's response, it seems like the program understands human language and is an excellent writer. The program's response is grammatically correct and informative. ChatGPT can generate text about almost any topic. The program also requests user feedback so that it can improve responses. That a piece of software can "write" so well is truly astounding. Naturally, students (many of whom detest writing) have flocked to ChatGPT in vast numbers. One student at our university, when asked if she knew about ChatGPT, responded to us, "Oh, you mean the program we use to take online tests?" Consequently, teachers, lecturers, and professors are having to develop strategies for "coping"

with and identifying students who use ChatGPT without authorisation (Greenlaw et al., 2023). ChatGPT does suffer from misinformation and hallucinations (Abreu, 2023), statements that are not based in fact but emerge randomly from the depths of the program. ChatGPT is not error-free. Obviously, such inaccuracies are highly problematic, and the program comes with a disclaimer.

It seems impossible to police student use of ChatGPT. In the program's own words, "Its availability as a 24/7 virtual tutor and quick responses makes it a valuable resource for students looking to enhance their learning experience. Additionally, ChatGPT can also aid in brainstorming ideas and refining essays or academic papers, further supporting students in their educational pursuits." Thus, few would argue about ChatGPT's utility. We choose to embrace it, so the research question we address is how ChatGPT can be incorporated effectively into a student's education. Our objectives is to come up with a set of guidelines to help instructors be better prepared to handle student issues involving generative writing tools. We employ our collective 75-plus years of experience in artificial intelligence, computing, writing, and teaching to formulate these guidelines. As authors of over 140 research papers and 35 books, including a book on technical writing, presenters of over 250 presentations, developers of accreditation criteria for ABET (the Accreditation Board for Engineering and Technology) and numerous programs in an assorted array of computing disciplines, we bring considerable expertise, depth of knowledge, and insights to this research.

Jarrah, Wardat, and Fildago believe that the use of ChatGPT by students in writing has created concerns and moral dilemmas (Jarrah et al., 2023). Bayne reveals that "Excessive dependence on AI may lead to inadequate comprehension of the material, resulting in unpreparedness for subsequent assignments" (Jarrah et al., 2023). Although ChatGPT may facilitate plagiarism and cheating, it can be used as a tool that improves students' writing skills without these misuses. Meyer et al., for example, clarifies that "A student that turns in verbatim text generated by ChatGPT in response to a single request from that student to write that essay would be clearly undesirable; in contrast, a student who produces an essay by taking on the roles of prompt engineer, fact checker, and editor, should be viewed positively" (Meyer et al., 2023). Despite the benefits that ChatGPT has, certain strategies need to be put in place to avoid the misapplications of ChatGPT, such as copying text verbatim from ChatGPT for answering essay assignments. According to Rahman and Watanabe, one of the strategies to prevent abuse of ChatGPT is "an advanced plagiarism-detection tool that can be used to detect AI-generated texts" (Rahman and Watanobe, 2023).

Having already given a thorough literature review in this introduction, the remainder of this paper is organised as follows. In Section 2, we describe the issues involved with students using ChatGPT and the new writing skills that are required to get the most out of the program. We present a brief high-level description of how ChatGPT works in Section 3. Bringing to bear our experience as authors, producing camera-ready papers and books in dozens of instances, and our work as professional editors, in Section 4, we discuss developing editing skills. Relying on our 45-plus years of experience in accreditation, among other things, in Section 5, we present material on fact-checking skills. In Section 6, we discuss how to phrase queries to ChatGPT in order to produce better output. In Section 7, we discuss abuses and disguises of ChatGPT's writing. In Section 8, we present a summary and conclusions.

CHATGPT Writes Better Than Our Students

In an earlier paper (Greenlaw et al., 2023), we examined how well ChatGPT writes versus computing students from three university classes: a first-year programming fundamentals course, a second-year computer organisation and architecture course, and a fourth-year entrepreneurship and management of IT systems course. Based on numerous quantitative metrics such as grammatical correctness, spelling errors, average word length, vocabulary, sentence structure, word repetition, and punctuation usage, we found that ChatGPT writes considerably better than our average university computing student. We are generalising here, but in discussions with colleagues from other departments, they draw similar conclusions. Table 1 summarises our findings. We also learned that students are not able to make any significant improvements to ChatGPT's output (Greenlaw et al., 2023). That is, they either do not have the editing skills or they were hesitant to improve ChatGPT's writing. We found that few students could identify ChatGPT's hallucinations and would be better served with improved fact-checking skills.

Table 1. ChatGPT versus student writing (Greenlaw et al., 2023).

Metric	Comparison	Better Performer	Comments
Subjective quality of the writing as judged by the professor	ChatGPT performed 10% better.	ChatGPT	The overall structure, organization, content, and writing were one full grade level better by ChatGPT.
MS WORD's grammatical correctness	ChatGPT had seven times fewer errors.	ChatGPT	ChatGPT made far fewer grammatical mistakes than students.
Accuracy of information	ChatGPT had two more errors.	Students	ChatGPT errors were hallucinations. Such factual errors can be misleading and dangerous.
Spelling errors	Student writing contained one more error per page.	ChatGPT	ChatGPT did not make spelling errors.
Word length	ChatGPT words were 20% longer on average.	ChatGPT	ChatGPT used longer words, and it used more descriptive adjectives.
Vocabulary	ChatGPT's vocabulary was 5% larger.	ChatGPT	ChatGPT used a larger vocabulary, which helped to improve readability.
Sentence structure	ChatGPT used a 20% wide range of sentence structure.	ChatGPT	ChatGPT used a greater number of sentence styles. This variety tended to make the writing more interesting to read.
Word repetition	The writings were similar.	Tie	There were no significant differences found via this metric.
Punctuation mark usage	ChatGPT used 20% more punctuation.	ChatGPT	ChatGPT used significantly more punctuation, perhaps due to the greater variety of its sentence structures, but also

			attributable to having more information, and therefore listing a larger set of items when relevant.
--	--	--	---

In curricula that we are familiar with at all levels in all disciplines, instructors do not focus on writing components of their courses on editing and fact-checking. In Greenlaw's book on technical writing (Greenlaw, 2012), there is a discussion about editing and copyediting, but the phrase "fact checking" is not mentioned. Although assumed, fact-checking is simply not a skill that was needed in the context of writing in the past. The authors did their research in reliable sources and then cited those facts. There was a limited need to verify information from these trusted sources. Of course, a second or third source might have been consulted, but fact-checking in the sense that we are describing here was not necessary or performed. Published material was thoroughly reviewed, copyedited, and for the most part, accurate. When Wikipedia became popular at the beginning of the 21st century, authors were hesitant to cite the online encyclopedia. However, over time, the website has become better known and is often cited, or at least used as a starting point for an author uncovering original sources.

With millions and millions of students using ChatGPT, there needs to be a shift in the way writing is taught. Purists may feel otherwise, but any open-minded reviewer is likely to reach the same conclusions as reported in (Greenlaw et al., 2023). ChatGPT's output, as measured by analytical metrics or via subjective means, is superior to that of most university students (Greenlaw et al., 2023). Sceptics may claim otherwise, but no one can help but notice that AI generative writing technology is rapidly improving. We anticipate a widening gap in the quality of writing done by AI tools versus humans. Will the need for people to write well disappear? Absolutely not. However, we believe the AI tools will soon be producing higher-quality writing than all but the best writers.

There is much writing that needs to be done where a large language model's utility is limited, for example, when writing timely memos, instructions, material involving proprietary, personal, or classified information, or breaking news. In other words, a program such as ChatGPT cannot write about something for which it does not yet have relevant data. It could help with the structure, though or other generic parts of the writing. As a concrete example, consider this paper's authors' intercommunication. When individuals are collaborating, they need to indicate the status of an item, what remains to be done, how that work is to be divided, what the timeline is for completion, and so forth. As of this writing, it seems impossible that ChatGPT could write an email from scratch containing such information on behalf of one of the authors. A snippet of which might read as follows:

"I have completed about 85% of the paper. Please fill in items that I have highlighted in yellow. Include information for your classes in Tables 1 and 2, and please fill in the required data on page 10. Also, add a literature review and citations. I have highlighted in yellow where that material belongs. When we meet next Wednesday, I can answer any questions and explain my methodology."

The only person with the "data" to write such an email is the author doing the actual writing. This information cannot be mined from the WWW using current technology. Thus, individual writing skills are going to be needed for the foreseeable future. We believe that ChatGPT and other generative AI writing tools can be used to help develop a student's writing ability. Of course, if we have no humans generating well-written material, where is ChatGPT going to get its training data? As discussed in Section 3, ChatGPT needs an incredible amount of data (written material) to train its artificial neural network. The higher the quality and the more accurate the data are, the better ChatGPT can perform.

In addition to editing and fact-checking, a third skill that students need to develop is how to phrase input to ChatGPT effectively (also known as prompt engineering). For the moment, this is a bit of an art form. That is to do things well the first time a user enters a prompt to ChatGPT. For others, it is a process of trial-and-error. Enter a prompt; see what output ChatGPT produces; intelligently revise the prompt to guide ChatGPT in the desired direction. Later in this paper, we include a discussion about phrasing queries to ChatGPT in an effective manner.

Essence Of How CHATGPT Works

Having at least an intuitive sense of how ChatGPT works is useful for understanding the writing that the program generates in terms of accuracy of information, word usage, grammar, and style. This potentially can help one identify text written by ChatGPT, as opposed to a student. In this section, we discuss the inner workings of ChatGPT at an intuitive level. Stephen Wolfram's excellent exposition (Wolfram, 2023) inspired this discussion. Our goal is not to explain the full technical details of ChatGPT's architecture, but rather to give an idea of how the program works and of its enormous size. The program is extremely complex, and in a precise mathematical sense, no one fully understands exactly why the program writes as well as it does (Wolfram, 2023). As mentioned earlier, OpenAI has released several versions of ChatGPT, including versions 1.5, 2.0, 3.0, 3.5, 4.0, 4o, 5.0, and 5 mini. For complete details of the various releases, see (ChatGPT, 2023). Sometimes these "improved" versions perform worse than earlier versions (Duarte, 2023), at least in specific instances. The program is complex, as of this writing, having roughly 1.8 trillion "knobs" which need to be tuned, and not fully understood mathematically.

ChatGPT is a *large language model*. Its underlying architecture is an *artificial neural network* (Science Direct, 2023). The original artificial neural networks were designed to model how a human brain functions, having features for neurons and the pathways among them. In an artificial neural network, there is an input layer and an output layer. In between these layers sits a feed-forward only set of connections going from the input layer, passing through successive layers of neurons, the *hidden layers*, until the output layer is reached. The example artificial neural network shown in Figure 1 has four inputs, has three intermediate layers depicted with the ellipsis representing additional hidden layers, and has one neuron displayed in the output layer. The neuron receiving input i_3 has three pathways out, leading to all but the neuron shown under input i_2 .




Figure 1. An example of an artificial neural network having 4 inputs, 1 output, 12 “visible” neurons, and three intermediate layers explicitly shown. The ellipsis represents additional hidden layers.

According to a query to ChatGPT itself, its artificial neural network has 96 layers, but the number of layers may vary from version to version. Each neuron is like a gate in a circuit. However, at a neuron, a general function of its inputs is computed rather than, say, just a simple Boolean function as in a gate in a Boolean circuit. Suppose a neuron labelled Z has inputs $\vec{x} = (x_1, x_2, \dots, x_k)$, where k is a natural number. Associated with the neuron Z are a set of real-valued weights $\vec{w} = (w_1, w_2, \dots, w_k)$. At the neuron Z , the value

$$f(\vec{x} \cdot \vec{w} + b)$$

is computed, where ‘ \cdot ’ is the dot product, b is a constant value, and f is a non-linear activation or what is sometimes called a thresholding function. Thinking back to artificial neural networks trying to mimic the brain, the idea is for a neuron to fire if its input signals are activated by reaching a certain threshold. For each neuron in the artificial neural network, different sets of weights and different constants are typically used (Wolfram, 2023). In one version of ChatGPT 3, there are 175 billion parameters (OpenAI, 2023). As of this writing, it is estimated that ChatGPT may have as many as 10 trillion parameters. The activation function f is usually applied throughout the entire network. In other words, in practice, only one activation function is used. The b term is the *bias*. There is a magic to training the artificial neural network, through techniques such as *gradient descent* in *back propagation*, and getting the weights and constants adjusted properly (Wolfram, 2023). The same exact training data is often run through the network many times. One can think of adjusting thousands of billions of knobs to tune the artificial neural network. The number of parameters used in these large language models is growing rapidly and will soon number in the hundreds of trillions. Now, what role does the artificial neural network play in ChatGPT?

Although ChatGPT supports many languages other than English, we focus on English. Consider all the text available on the Web. According to Google in April 2023, there were over 30 billion indexed Web pages. Let us estimate that each Web page contains 1,000 words. This gives us about 30 trillion words of data. Note that the exact number is not important. The point is that the Web provides tremendous amounts of data. In fact, ChatGPT uses data from other sources as well. If there are 50,000 English words, there would be 7.7×10^{311} ways of “writing” a paragraph of just 100 words. Of course, many

of these paragraphs would not make any sense, and more words would mean many more possibilities. According to some sources, see (CK-12, 2023), for example, an upper-bound estimate on the number of atoms in the universe is 10^{82} . If we wanted to examine the contents of the entire Web combined with all other available data sources to predict “what 100-word paragraph might be written next,” we can see that there is not enough data available to do this task. That is, there are 7.7×10^{31} ways of writing such a paragraph, but there are only 300 billion such sample paragraphs on the WWW. In other words, each paragraph appears a vanishingly small number of times, if at all. Certainly not often enough to predict a “next paragraph.”

We can accurately estimate the probability of the next word to occur to complete a partial sentence; for that, we have enough data, but we cannot get a handle on what the probabilities of an entire “next paragraph” would be. The artificial neural network incorporated into the ChatGPT architecture can be thought of as supplying estimates of the probabilities for such paragraphs (Wolfram, 2023). There is some randomness built into ChatGPT too, as the programmers at OpenAI realised that by always selecting the highest probability completion to write down next results in redundancy and “less interesting” and “boring” writing. Thus, in our writing experiments with students, we got different responses from ChatGPT for the same exact queries (Greenlaw et al., 2023). This is true in general. If you enter the same prompt to ChatGPT, a different output will be produced each time.

In summary, we can think of ChatGPT as a huge, complex program based on an unfathomable amount of data that requires an incredible amount of computing resources, electricity, coolant, and money to run, that in essence “predicts” what to “write” based on what it has learned about language through its artificial neural network. If the data used to train a large language model is flawed, then so will its output be incorrect. And, of course, much of the “information” on the Web is inaccurate, for example, fake news, opinion pieces, and a proliferation of existing errors. Filtering and feedback mechanisms, giving users the chance to rate ChatGPT’s output, can be used to improve ChatGPT’s accuracy, at least in theory. But, for the time being, ChatGPT does not understand language in the same way a human does. It takes the bulk of what has ever been written in English in the form of Web content, books, articles, and so forth, and outputs a random high-probability text based on a user’s input. Even if only accurate information is used as data, inaccurate results can be produced due to the manner in which large language models are trained. Hallucinations aside, ChatGPT does such a remarkable job is truly amazing and a remarkable technological feat! One that no one envisioned as few as a couple of years ago.

Editing-Skills Development

Traditional undergraduate courses in most fields include writing assignments. Accrediting agencies often require students to be proficient in written and oral communication (ABET, 2023). Namibian Universities such as the University of Namibia (UNAM), the Namibia University of Science and Technology (NUST), and the International University of Management (IUM) include compulsory language courses for first-year students aimed at developing written and oral communication skills. In our experience, the writing assignments in such modules do not focus on editing. Rather, they focus on general writing skills. Even with the introduction of ChatGPT, the need to teach general writing skills remains essential, but assessing this ability may become more difficult for instructors.

In some respects, editing requires a greater attention to detail than writing. Copyeditors are typically people with an excellent eye for detail. Although revising and editing might be part of a unit for a writing course at a university, much emphasis was traditionally placed on the writing itself rather than editing. Lecturers emphasised teaching students how to write rather than how to edit their written work afterwards. Students are unlikely to excel in their academic writing skills if they are not made aware of the errors in their writing and learn how to rectify them. With AI generative writing tools freely available that are producing quality output, we believe this calls for a shift from a focus on writing to editing skills. Many students have discovered software editing tools to assist them in writing. One of the common ones mentioned earlier in this paper is Grammarly (Grammarly, 2023). If an author is not careful and knowledgeable, such software might mislead a writer and introduce errors. In running the final version of this paper through WORD's grammar reviewer, seven grammar "issues" remained. However, none of them constitutes a grammatical error. The "errors" reported by WORD reveal the grammar checker's limitations. It does not actually understand language.

Students need a command of basic language skills to be able to identify errors in writing and to improve an existing piece of writing. Some of the basic grammatical skills are tenses, punctuation, capitalisation, article usage, pronoun selection, and so forth. Table 2 provides a list of editing skills that can help students improve a piece of writing, specifically a piece of writing that is computer-generated. Column Three of the table contains comments about how well ChatGPT wrote versus students in comparison of their writings, and based on our experience in analysing and grading writing assignments.

Table 2. Mental checklist for editing skills.

Skill or Item to Consider	Purpose	Comments about ChatGPT
Consistency	Have the writing flow smoothly and make logical sense. Ensure that writing maintains a consistent tone, style, conventions, and tense throughout.	The overall structure, organisation, content, and writing were one full grade level higher in ChatGPT's output than in student writing.
Grammatical correctness	Remove grammatical errors and improve the clarity of the writing.	ChatGPT made far fewer or no grammatical mistakes common in students' writing.
Spelling	Remove spelling errors.	ChatGPT did not make spelling errors.
Tenses	Use accurate and consistent tenses.	For the usual simple tenses, ChatGPT rarely makes errors.
Vocabulary	Accurate conveyance of meaning; make the writing interesting and read well; allow for the audience's comprehension.	ChatGPT used greater length words, and in general, more descriptive adjectives. ChatGPT used a larger vocabulary which helped to improve readability.
Sentence structure	Add variety. Make the writing interesting.	ChatGPT used a greater number of sentence styles which tended to make the writing more interesting to read and less redundant.
Punctuation	Make the exposition clear and provide an easy-to-follow style. Improve the flow of the writing.	ChatGPT used considerably more punctuation than students.
Cohesion	Link ideas together. Improve the flow of the writing.	ChatGPT incorporated many more ideas into its writing.

Fact-Checking Skills Development

ChatGPT has the ability to write well and incorporate factual information, especially using data that is publicly and legally available. Its ability to provide real-time information, personal opinions and experiences, private and confidential information, and live updates on contemporary issues seems to be limited, at least as of this writing. Even with the “factual information” generated by ChatGPT, however, it is essential that an author thoroughly processes all the information, including items that appear to be facts or common knowledge. An author must look up and confirm every generated detail to confirm its validity and verify cited sources to ensure they actually exist. As described in section 3, the manner in which ChatGPT produces its output, by “analysing” huge volumes of data and “piecing” that information together via its artificial neural network, does not lend itself well to citations, although this aspect of large language models is improving rapidly. As of this writing, when ChatGPT gets something wrong and makes a hallucination, there is no telling exactly why or how an error was made. The program does not simply regurgitate information and proliferate existing errors. Its processing of information through its immense artificial neural network is far more complex than that, as outlined in section 3. ChatGPT itself had this to say about its responses:

“My responses are based on the patterns and information present in the data on which I was trained. If the training data contains inaccuracies or biases, it could potentially affect the accuracy of my responses. However, I don’t have the ability to independently verify the accuracy of the information, and my responses should be taken as generated outputs rather than absolute truths. It’s always a good practice to cross-check information from multiple sources, especially for critical or sensitive topics.”

As described in section 3, ChatGPT is based on an enormous artificial neural network which is trained on a huge volume of data with perhaps tens of trillions of “knobs” being tweaked. When a thought emerges in a person’s mind, it seems impossible to “cite” exactly why or how that particular thought emerged. “Thought and consciousness are activated by sense-data, but on their turn initiate parallel and interactive processes in the brain where both lower levels (e.g., perception) and higher levels of cognition (e.g., memory, conceptual systems, and world knowledge) occur simultaneously and interactively” (Rahman and Watanobe, 2023). For a given sentence output by ChatGPT, there is no way to know directly where its wording came from, nor to cite sources in the traditional way, although new technology is being developed which attempts to perform this task. ChatGPT’s output is, in a sense, a conglomeration of all its training data. A major open problem is how a large language model can verify the accuracy of its output and eliminate hallucinations. If one thinks of simply generating the next word in the output based on what has been generated previously and on mathematical probabilities, then all instances of those words in the input data impact the choice of the next word. There is no way to “cite” where the “information” originated. The model underlying ChatGPT is, of course, far more complex and intertwined than what the previous sentence assumed, as sketched in section 3. Because all the input data impacts the output, there is no way to cite “information,” at least not in the traditional manner.

For a person using ChatGPT in any serious way, there is a need to verify and corroborate the validity of the program’s output by checking multiple sources. Due to the complexity, randomness, and the magic behind the way ChatGPT works, which is not fully understood mathematically by anyone, and not being able to cite sources in the usual sense directly from where the program gets its “facts,” we concluded that it is necessary for educational institutions to modify curricula so students can use

generative AI programs more effectively. This shift involves teaching students new skills, which traditionally were not taught, as well as emphasising and expanding the coverage of other topics such as library-science citations. Table 3 provides a list of fact-checking skills and qualities that we consider essential for students to verify the accuracy of writing generated from large language models. They need to do detective work and learn to question information as to its reliability. Incorrect information needs to be edited or deleted.

Table 3. Fact-checking skills that are critical to verifying that a generative AI program's output is correct.

Skill or Item to Consider	Purpose	Comments about ChatGPT
Critical questioning	To ask relevant questions about information provided. To doubt information until it is verified.	Learn to formulate and pose questions such as the following: What are corroborating sources? What evidence is provided? Are there possible biases? Is a given statement true?
Question the validity of every statement	Uncover errors of fact.	Find reliable sources that substantiate or refute statements.
Verify numbers	Ensure stated values, quantities, and numbers are correct.	Confirm that all values are accurate through reliable sources.
Do detective work	Make sure statements are related and follow logically. Find problems in the writing.	Uncover false statements. Delete and remove unneeded text. Improve the writing.
Verify names	Improve the accuracy of the writing.	Check that all names are correct and represent real people and places. Verify that a person actually did what was attributed to them. Verify locations of places. Verify titles. Verify sources.
Identify misinformation	Use multiple reliable and reputable sources to verify information.	Check that all information is correct, regardless of how plausible and official it might seem. Be suspicious of compound information, as only part of it may be correct.
Ask questions	Get in the habit of doubting everything.	Learn to pose questions such as: What could be wrong with this? Does this agree with what I know as factual? Where can I corroborate this information? Does this statement make sense?
Search carefully and thorough	Employ prompt engineering techniques to ask better questions.	Pose additional questions and see if contradictory information is produced. Use reliable search engines and sources to verify facts.
Check for consistency	Verify that information follows logically and there are no contradictions.	Use multiple, reputable sources and make sure they agree with the item being checked.
Analyze information	Assess the plausibility of information and determine if it requires further investigation.	Do this to achieve consistency and accuracy. Never assume something is valid. Check everything, even small trivial items.
Pay attention to details	Verify obvious information and check for deviations such as dates,	Check even the most innocuous-looking trivial items for their veracity, as any errors

	actors, authors, and so forth.	can bring into question the validity of an entire document.
--	--------------------------------	---

Table 3 is not intended to be comprehensive but provides a substantial list of skills that can help students debug, verify, cite, and edit output from a generative AI program. Rather than teaching traditional writing skills, we believe these items should become more of a focus. Throughout a curriculum, class time that was spent on developing writing skills can be used to teach the items in Table 3. We believe the result will be accurate and improved “student” writing.

How To Query CHATGPT Effectively

As discussed in section 3, the same query entered repeatedly to ChatGPT will produce different outputs. So, what is the most effective type of input to ChatGPT? Our experience has shown that the following items are important to consider:

Provide ChatGPT with context, detailing specific facts as much as possible. For example, “provide a detailed overview of the historical development of artificial intelligence, highlighting key milestones, breakthroughs, and the evolution of AI technologies” instead of simply “give me an overview of the development of artificial intelligence.” Although both inputs will generate a reasonable response of the historical development of AI, the first will include a specific focus on the milestones and breakthroughs that mark the historical development of AI.

Be as precise as possible with the input and guide ChatGPT as to the information that is important or of interest. For example, “provide a one-page biography of Jasmin Paris, the female long-distance runner who lives south of Edinburgh, Scotland and became the first woman to complete the gruelling Barkley Marathon in Frozen Head State Park in Morgan County, Tennessee in 2024” instead of “who is Jasmin Paris.” It may help to add information about the person doing the prompting as well, for example, prefixing the above prompt with “I have run 50 ultramarathons” may lead to a more informed response. ChatGPT can use its knowledge of the prompter to provide a more suitable answer to the prompt.

Be direct and unambiguous. For example, “write three paragraphs about Marc Andreessen, who programmed the first graphical web browser called Mosaic and founded a company called Netscape Communications Corporation” instead of “write about a programmer named Marc Andreeseen.”

Strive to use keywords that have a single meaning. For example, the word ‘panthera onca’ is the scientific name for the large cat species known as ‘jaguar’, instead of ‘jaguar’, which is the name of an expensive, luxury car; a large spotted cat found in the Americas; and a professional American football team.

The examples provided give a flavour of the manner in which prompts should be designed and entered into ChatGPT. Notice we used the word ‘designed’, although some texts refer to it as engineering. Careful thought needs to go into ChatGPT’s input. We have found empirically that more specific, targeted prompts can, in some cases, reduce the number of hallucinations that ChatGPT produces.

Abuses And Disguises Of CHATGPT’s Writing

As noted earlier, when asked about their familiarity with ChatGPT, one of our students responded, “I use ChatGPT for online exams.” Students are using ChatGPT in unexpected and unethical ways, and

those ways will multiply and expand in unpredictable directions. This is not too surprising, as the program burst onto the scene and is not well understood by the academic community. Guidelines for usage are emerging, but not as quickly as needed. As of this writing, tools to police the use of the program are not very effective. Among other things, students use ChatGPT to complete programming assignments, write term papers, generate figures and tables, solve problems, develop essays, conduct research, and produce presentations. Students have even developed techniques to disguise their use of ChatGPT. We list some of these insidious methods next.

How to fool teachers by editing ChatGPT's output:

- Introduce misspelt words.
- Introduce minor, simple errors of fact.
- Reduce the amount of punctuation used.
- Delete flowery adjectives.
- Replace words that one does not know the meaning of, thereby using a reduced vocabulary.
- Replace longer words with shorter, simpler synonyms or a phrase.
- Add in some redundancy or duplicate words to introduce a grammatical error.
- Transpose a pair of words so they are out of order.
- Break long sentences into simple sentences.
- Introduce grammatical errors.
- Disrupt the logical flow of something, for example, swap the order of a couple of sentences.

As students become more adept at fooling teachers, they also need to improve their detective skills, looking for student-inserted material intended to throw them off the trail. Although most students in our experience have good ethical standards, some will not follow the rules. To cope with these students, teachers need to develop new skills and gain a greater understanding of ChatGPT and its uses.

Summary and Conclusions

AI generative writing tools are a truly extraordinary advance in technology. Not only do they produce exceptional quality output, but they also fill a need. That is, most students do not like to write and perform research. ChatGPT reduces these two time-consuming tasks to trivial and immediate tasks. Regardless of guidelines developed to restrict student use of AI generative writing tools, students are going to use them. These tools are writing better than our students. They seem unable to improve ChatGPT's output other than to correct hallucinations. This suggests that students should be taught editing, fact-checking, and prompt-engineering skills with a reduced emphasis placed on writing skills themselves. Considering this point, it is recommended that a course/portion of a course addressing these three skills be implemented for first-year university students in order to make them aware of

the issues involved in using AI generative writing tools, and thereby, making them better writers. Besides, when students become better writers, we believe they get less tempted to cheat, copy, or commit any irregularities from ChatGPT. Furthermore, they will start seeing ChatGPT as a great and useful tool to improve their writing skills rather than a pillar to lean on (copy and paste). Editing and fact-checking are not skills that have been typically stressed or even taught in computing curricula. And, of course, prompt engineering, as used in this paper, was never taught.

The authors have first-hand experience with a number of trends in education. When the calculator first came out in the early 1970s, replacing the slide rule, many educators were reluctant to let students use calculators for exams. Gradually, calculators became more accepted, and less time was spent in curricula teaching simple mathematics. Handwriting courses were replaced by typing courses. As computers became ubiquitous in many institutions, computer literacy courses became compulsory for all students. (One still is, at our institution.) Many institutions are introducing a general course on cybersecurity, so that members of society know how to protect themselves online and use technology without fear of being hacked. Somewhat analogously to these advances and shifts, we are suggesting a move to editing, fact-checking, and prompt engineering skills. We have provided many guidelines and suggestions in this paper as to what specifically might be taught. More research needs to be conducted along these lines.

As generative AI writing tools flood into educational institutions, teachers at all levels, from primary schools to post-graduate schools, are going to need additional training and education. It behoves such institutions to provide opportunities for their employees to get this much-needed background. Just as many publications have needed to rewrite their submission guidelines to address AI writing tools, institutions need to develop clear policies as to what is expected of students and instructors. Such policies must be enforceable. Instructors need a way to differentiate student writing from writing produced by ChatGPT, and automating that process is a good direction for further research.

References

- ABET, 2023. ABET. [Online] Available at: <https://www.abet.org>. [Accessed 27 September 2023].
- Abreu, J. S., 2023. Founding Fathers of Artificial Intelligence. [Online]. Available at: <https://quidgest.com/en/blog-en/ai-founding-fathers/>.
- CERN, 2023. A Short History of the Web. [Online]. Available at: <https://home.web.cern.ch/science/computing/birth-web/short-history-web>
- ChatGPT 3.5, 2023. ChatGPT 3.5. [Online]. Available at: <https://chat.openai.com/>. [Accessed 6 September 2023].
- ChatGPT, 2023. ChatGPT—Release Notes. [Online]. Available at: <https://help.openai.com/en/articles/6825453-chatgpt-release-notes>. [Accessed 12 September 2023].
- CK-12, 2023. How many atoms exist in the universe? [Online]. Available at: <https://www.ck12.org/flexi/earth-science/atoms-and-molecules/how-many-atoms-exist-in-the-universe>. [Accessed 12 September 2023].
- Davis, Z., 1986. Writer's Workbench. PC Magazine, May, pp. 190-199.

- Duarte, F., 2023. Number of ChatGPT Users. [Online]. Available at: <https://explodingtopics.com/blog/chatgpt-users>. [Accessed 6 September 2023].
- Gorin, R., 2009. Ralph Gorin, talk, gold medal for spell checker. [Online]. Available at: <https://exhibits.stanford.edu/ai/catalog/mz021fp0267>. [Accessed 5 September 2023].
- Grammarly, 2023. Grammarly. [Online]. Available at: <https://www.grammarly.com/>.
- Greenlaw, R., 2012. Technical Writing, Presentation Skills, and Online Communication: Professional Tools and Insights. xxi + 225 pages: IGI Global.
- Greenlaw, R., Mufeti, K., and Mwatilifange, S., 2023. Experiences with Artificial Intelligence Generative Writing Tools in Computing Curricula, 18th Annual International Conference of Education, Research and Innovation, Seville, Spain, November 2025. To appear.
- Horn, B. and Winston, P., 1975. Personal Computers. *Datamation*, 21(5), pp. 111-115.
- Jarrah, A., Wardat, Y., and Fildago, P., 2023. Using ChatGPT in academic writing is (not) a form of plagiarism: What does the literature say?. *Online Journal of Communication and Media Technologies*, 13(4).
- Meyer, J. G. et al., 2023. ChatGPT and Large Language Models in Academia: Opportunities and Challenges. [Online]. *Biodata Mining* 16(20). Available at: <https://biodatamining.biomedcentral.com/articles/10.1186/s13040-023-00339-9>. [Accessed 21 October 2025].
- Moore, G. E., 1965. Cramming More Components onto Integrated Circuits. *Electronics*, pp. 114-117.
- Moor, J., 2006. The Dartmouth College Artificial Intelligence Conference: The Next Fifty Years. *AI Magazine*, 27(4), pp. 87-89.
- OpenAI, 2023. Models Featured in OpenAI Research. [Online]. Available at: <https://platform.openai.com/docs/model-index-for-researchers>. [Accessed 3 September 2023].
- Peterson, J. L., 1980. Computer Programs for Detecting and Correcting Spelling Errors. *Communications of the ACM*, 23(12), pp. 676-687.
- QuillBot, 2023. QuillBot. [Online]. Available at: <https://www.quillbot.com/>.
- Rahman, M. and Watanobe, Y., 2023. ChatGPT for Education and Research: Opportunities, Threats, and Strategies. *Applied Sciences*, 13(9).
- Science Direct, 2023. Artificial Neural Networks. [Online]. Available at: <https://www.sciencedirect.com/topics/computer-science/artificial-neural-networks>. [Accessed 12 September 2023].
- Wikipedia, 2023. Spell checker. *Wikipedia: The Free Encyclopedia*. [Online]. Available at: https://en.wikipedia.org/wiki/Spell_checker.
- Wolfram, S., 2023. What is ChatGPT Doing ... and Why Does It Work? [Online]. Available at: <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/>. [Accessed 27 July 2023].