

Embracing Artificial Intelligence to Improve Self-Directed Learning: A Cybersecurity Classroom Study

Jim Marquardson
jimarqua@nmu.edu
College of Business
Northern Michigan University
Marquette, MI 49855, USA

Abstract

Generative artificial intelligence (AI) tools were met with a mix of enthusiasm, skepticism, and fear. AI adoption soared as people discovered compelling use cases--developers wrote code, realtors generated narratives for their websites, students wrote essays, and much more. Calls for caution attempted to temper AI enthusiasm. Experts highlighted inaccuracies in some of the generated information. Artists, writers, programmers, and other professionals who create for a living questioned their job security. Educators worried that these tools would harm academic integrity. Nevertheless, the tools persist, and AI enthusiasm shows no sign of abating. The current study focuses on how AI tools can be leveraged to improve self-directed learning in the classroom. In a capstone cybersecurity course, the instructor encouraged students to use ChatGPT 3.5 on a self-directed learning assignment to select topics, find learning resources, and carry out their learning plans. Some students did not use ChatGPT. Those who used ChatGPT found it to be a helpful learning aid. All students reported that they plan to use ChatGPT for self-directed learning after graduating. No student violated academic integrity policies. The results demonstrate that AI tools can enhance self-directed learning, though students should be trained to use them effectively. Recommendations for developing assignments that support academic integrity and rigor in the age of AI are provided.

Keywords: artificial intelligence, self-directed learning, cybersecurity, pedagogy

Recommended Citation: Marquardson, J., (2024). Embracing Artificial Intelligence to Improve Self Directed Learning: A Cybersecurity Classroom Study. *Information Systems Education Journal*, 22(1), pp.4-13. <https://doi.org/10.62273/WZBY3952>

Embracing Artificial Intelligence to Improve Self-Directed Learning: A Cybersecurity Classroom Study

1. INTRODUCTION

ChatGPT took the world by storm in 2022. Just two months after its launch, ChatGPT achieved 100 million monthly active users (Hu, 2023). The impacts of AI on business, political discourse, economics, and other important areas remain unclear, but AI tools already demonstrate their disruptive potential. Educators must prepare students to embrace change in a world where AI tools become ubiquitous.

Companies are already adjusting their workforce planning to account for AI. For example, IBM plans to pause hiring for roles that AI can replace (Ford, 2023). While many roles cannot be replaced entirely by AI, companies may increase productivity expectations for existing employees (Reddington, 2023). Professionals can position themselves for success by embracing change and honing new AI skills. Therefore, educators should teach the appropriate use of AI tools. People need to learn to use AI tools just like they would any other advanced technology. Academia has an opportunity to prepare students to succeed in the changing world where AI tool use is the norm.

Concerns about academic integrity and rigor complicate AI tool adoption in academia. Using AI, students can find quiz answers, generate essays, and otherwise circumvent assessment of learning. Students may attempt to offload critical thinking to AI tools. Educators must find ways to enhance learning while maintaining academic integrity.

This paper addresses the use of AI in academia by providing a brief overview of AI history and tools, then presents a study of adopting AI in a self-directed learning assignment. Recommendations for adopting AI in the classroom are provided.

2. LITERATURE REVIEW

This section provides an overview of generative artificial intelligence terms, tools, and techniques relevant to the current study. The aim is not to provide an exhaustive explanation of AI, but to give a brief historical context relevant to current AI tools and discuss emerging concerns in academia.

Artificial Intelligence

In 1955, John McCarthy and colleagues coined the term “artificial intelligence” in a research project proposal stating that, “The study is to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it” (McCarthy et al., 1955, p. 1). Defining and creating artificial intelligence proved more challenging than first imagined. Advances in hardware and software yielded tremendous gains in computing power, yet for decades true artificial intelligence felt out of reach. Enthusiasm and funding for AI research waned in the 1970s and 1980s—a period known as the “AI winter” (Haenlein & Kaplan, 2019).

In the 1980s and 1990s, researchers continued to make progress building intelligent systems. For example, in 1997, IBM’s Deep Blue system famously beat the world chess champion Gary Kasparov at chess (Campbell et al., 2002). But perhaps because Deep Blue could only play chess and its programmatic logic could be explained in a relatively straightforward way, questions lingered as to whether Deep Blue was truly a manifestation of artificial intelligence.

In subsequent decades, data mining and machine learning demonstrated the power of statistical inference on large datasets. These techniques gained traction in academia and industry. Deep learning is a branch of machine learning that extracts meaningful features from data to construct artificial neural networks (LeCun et al., 2015). These artificial neural networks are then used to make predictions based on new inputs. Researchers used deep learning to drive advances in image recognition, speech recognition, and recommendation systems. Deep learning proved the value of training models on vast quantities of data to improve predictive performance. The models are often evaluated by their accuracy. For example, the percentage of pictures of cats were correctly classified by the machine learning algorithms as cats.

Large Language Models

Large language models (LLMs) are AI tools used to understand and generate human language. In 2017, researchers introduced transformers which

allowed LLMs to better capture contextual information than previous techniques (Vaswani et al., 2017). Models using transformers benefit from parallel processing. This architecture improvement coupled with GPU-accelerated computing vastly reduced the cost of training LLMs on large datasets. The "GPT" in ChatGPT stands for "Generative Pre-trained Transformer."

LLMs are composed of parameters—connections between artificial nodes in a neural network. A **model's capacity to learn increases as the number of parameters increases**. The number of **parameters from ChatGPT's underlying models** increased from 175 billion in GPT-3 (Brown et al., 2020) to 1.4 trillion in GPT-4 (Katz et al., 2023).

Generative AI

Generative AI systems create new content. Examples of content that generative AI currently creates include poems, videos, study guides, music, and source code (Longoni et al., 2022). Generally, a person gives the generative AI tool a prompt, and the AI tool generates the output. With some tools like ChatGPT, context is retained between prompts until a new chat is started.

Copyright Concerns

Developers train models on vast amounts of data. Some of the data may be copyrighted, and it is currently unclear how intellectual property laws apply to AI-generated work. Some artists lament that AI tools improperly mimic their unique artistic styles (Chayka, 2023). Stability AI trained its Stable Diffusion AI image generation system on data that included copyrighted images from Getty Images (Brittain, 2023). Some of the novel images produced by Stable Diffusion based on **user prompts included Getty's watermark**—an indication that the output was not merely inspired by existing work, but actually incorporates it in the output.

Information Accuracy

AI tools sometimes provide misinformation. Below are some personal interactions with tools such as ChatGPT, Google Bard, and Bing Chat that yielded unexpected results.

- Two of three chat AI tools incorrectly identified the second layer of the OSI **model as the "network" layer**.
- ChatGPT generated code to perform XOR decryption with a specific key and ciphertext. ChatGPT also included the result of running the supplied code. The code worked correctly when ran locally, but the plaintext it said it produced from running the code was not correct.

- ChatGPT gave examples of PowerShell code to automate Windows Server administration tasks. Some of the PowerShell modules it recommended do not exist.

AI image generators sometimes produce people with the wrong number of hands, impossible building architecture, and other visual artefacts that ruin the illusion. Tools that produce text sometimes invent information, such as research article titles that do not exist. These incorrect or unexpected outputs are often labelled "**hallucinations**" (Zhang et al., 2023). However, **the term "confabulation" may be more appropriate**. The neuropsychological literature describes a confabulation as the unintentional production of false memories without the intent to deceive (Kopelman, 1987).

AI tools currently occupy an interesting space in which they sometimes show brilliance, but also create monumental factual mistakes. For example, ChatGPT (powered by the GPT-4 model) earned a passing grade on the Uniform Bar Exam (Katz et al., 2023). Despite its apparent **legal prowess, relying on ChatGPT's output landed one lawyer in hot water**. In 2023, the lawyer relied on case law provided by ChatGPT to prepare a brief filed in federal court, but unfortunately, ChatGPT fabricated the case law (Weiser, 2023). The lawyer in question may face sanctions for failing to **critically evaluate ChatGPT's output**.

User Interfaces

Technically adept people have used machine learning and AI tools for years, but the hardware requirements, cumbersome software configuration, and use of the software proved challenging. Cloud-delivered AI tools offload the significant hardware requirements and configuration complexity.

For several years, companies have deployed chatbots to triage customer problems and provide more efficient customer service, so the idea of chatting with an artificial agent became a common experience. Today, chat-based tools provide an easy user interface for interacting with **AI models. ChatGPT's web interface is simple**: start a new chat and send a message. Chat prompts can be reframed and retried. Midjourney leverages the group chat application Discord for its user interface. Users enter prompts in the group chat and wait for the server to generate their images.

More technically sophisticated users can access AI services using application programming

interfaces (APIs), but many people find value in the chat interfaces alone.

Adoption

Popular theories help shed light on the rapid adoption of AI tools. The Unified Theory of Acceptance and Use of Technology (UTAUT) explains that performance expectancy (i.e., the predicted benefits of using technology), and effort expectancy (i.e., the ease of use) predict adoption (Venkatesh et al., 2012). Anecdotal reports support ChatGPT's ease of use (e.g., Marcus, 2022) and usefulness.

Task-technology Fit (TTF) theory is another useful lens for explaining generative AI's strong adoption. The theory proposes that people adopt technology when task characteristics and technology characteristics align (Goodhue & Thompson, 1995). Examples of tight alignment between real-world tasks and generative AI capabilities abound. For example, programmers need to write code to meet specifications, and ChatGPT can write code (Israelsen, 2023). Real estate agents must create home descriptions, and ChatGPT can generate them based on prompts (Kelly, 2023). People often need to create text, art, code, and other deliverables. Generative AI tools have proven that their technological characteristics align with the task characteristics that people need.

Ethical Considerations

AI users may harbor apprehensions about data privacy, consent, and biases. Prompts, files, and any other data uploaded to AI tools could be tracked to an individual. People may be hesitant to use AI tools for sensitive topics for fear of data leaks. Also, AI model bias might favor different groups which could reinforce stereotypes. People should have the opportunity to provide informed consent when using AI tools. The average consumer can choose to use AI, but in a classroom setting, it can be debated whether a student could provide meaningful consent if required to use AI tools as part of learning exercises. The following section delves deeper into AI adoption concerns in academia.

3. AI ADOPTION CONCERNS IN ACADEMIA

Options of AI tool use in academia are mixed. Some educators advocate for the prohibition of AI tools in academia. Others take an optimistic view of AI tool adoption in academia, arguing that, "with the right approach, [ChatGPT] can be an effective teaching tool" (Roose, 2023, para. 18). This section discusses some of the major

concerns and approaches for using AI tools.

Information Accuracy

Harvard University's motto succinctly summarizes the aims of academia—"veritas" meaning "truth." Knowledgeable experts can identify misinformation in AI tool output, but a novice may overly rely on the tools. Educators must teach students appropriate skepticism in the information they consume, especially in the case of current AI tools. Educators should encourage students to verify information, use multiple sources, and think critically when consuming information.

Academic Integrity

Educators cite cheating as a top concern with AI adoption (Blöse, 2023). New York City schools initially banned the use of ChatGPT (Rosenzweig-Ziff, 2023), but within months those bans were lifted (Klein, 2023). Because generative AI writes compelling text, it can be difficult to discern whether a student wrote an essay or simply asked ChatGPT to create it. Attempts to catch cheaters have proved challenging. One professor copied student essays into ChatGPT, asked ChatGPT if it had written the essays, and accused many students of cheating (Verma, 2023). However, the version of ChatGPT the professor used lacked the ability to detect plagiarism and eventually the students were exonerated.

Assessment Considerations

If a single essay submission no longer demonstrates mastery of a topic, educators must find alternative assessment methods. Strategies for assessment in the age of AI include requiring submissions of drafts, education on plagiarism, and careful scrutiny of submitted work (Cotton et al., 2023). Other suggestions include oral assessments and applied problem solving through case studies (Malik et al., 2023).

Virtual Assistants

Students can use AI tools to augment course instruction. Instead of waiting for an instructor to answer an email, respond to a forum post, or provide classroom instruction, students can engage with AI tools on demand. One key benefit of AI tools is asynchronous communication (Cotton et al., 2023). Khan Academy, the popular online learning platform, introduced its AI-assisted tutor named Khanmigo which evaluates student work and finds errors in student reasoning (Fried, 2023). While Khanmigo boasts a slim feature set today, it demonstrates the ability for AI tools to intervene in real-time rather than waiting for students to ask for help.

In many cases, AI tool creation and adoption caught many educators flat footed. The rapid improvement in AI models and tools introduces new challenges and opportunities. The current study provides data to inform the adoption and appropriate use of AI tools in academia.

4. METHODOLOGY

A study was conducted in a capstone cybersecurity course in a Midwestern university. One of the course learning objectives was to help students develop lifelong learning skills. This study replicates the classroom assignment more completely described in the article by Marquardson (2020). The major parts of the assignment are summarized in this section.

For the study, students completed a self-directed learning assignment. The assignment had four phases: the proposal, carrying out the learning plan, a presentation, and a reflection. For the proposal, students selected a topic, developed learning objectives, found resources to support their learning objectives, and documented how they would provide evidence of reaching those objectives. Once the instructor approved the plan, students carried out their proposals and learned their topic. Next, students presented a brief presentation of their topic to their peers. Finally, students submitted a reflection that included the learning evidence they promised in the proposal along with thoughts on the learning process.

The instructor encouraged, but did not require, students to use ChatGPT in each phase of the assignment. Students were told to cite ChatGPT in their initial learning plan and the follow-up learning reflections when they directly quoted text it generated. The instructor demonstrated use of ChatGPT in the classroom and highlighted examples where ChatGPT gave incorrect information.

After students submitted the final assignment reflection, they were invited to participate in the study with the promise of using the results to inform classroom discussions. Students did not earn course credit for participation or receive any other benefit. The study required that students complete a single survey with quantitative and qualitative questions. The full survey prompts are included in the results section. The survey included the option to upload ChatGPT logs. Identifying information was not collected unless students opted to upload their ChatGPT chat logs. The chat logs removed full anonymity because the filenames contained the student names, and

the topics in the log could be traced to that **student's unique topic**. In total, 9 participants (8 male, 1 preferred not to indicate gender) answered the survey questions, and 1 participant provided a sample ChatGPT interaction that can be found in Appendix A.

5. RESULTS

Of the 9 students who participated in the study, 7 elected to use ChatGPT. Summary statistics for those who used ChatGPT are included in Tables 1-3. Table 4 contains forward-looking statements about ChatGPT use, and therefore data from all 9 participants are included. The learning topics students chose included threat intelligence sourcing, python programming, Microsoft Active Directory, risk management frameworks, lock picking, web development, web server configuration, vulnerability scanning, JavaScript, and Docker.

Students reported their frequency of use on the self-directed learning assignment. Figure 1 shows the count of responses for all participants broken down by learning phase. Table 1 contains the summary statistics of those who reported using ChatGPT. The means for use in the three phases of the project fell between "Somewhat" and "Much."




Figure 1: Frequency of ChatGPT Use

How much did you use ChatGPT to help you:	Mean	SD
Select a Topic	3.4	1.3
Find Learning Resources	3.6	1.1
Learn Your Topic	3.4	1.0

Table 1: Frequency of ChatGPT Use, 1=Never, 5=A Great Deal, n=7

Because use of ChatGPT does not necessarily imply usefulness, the survey asked students to **assess ChatGPT’s helpfulness in completing the self-directed learning assignment**. The results in Table 2 indicate agreement that ChatGPT was helpful to some degree in all phases. Students reported that ChatGPT was a good teaching assistant, even favoring asking ChatGPT questions instead of their professor.

Prompt	Mean	SD
ChatGPT helped me select a relevant learning topic.	5.6	1.8
ChatGPT understood my learning goals.	6.0	0.8
ChatGPT helped me find good learning resources.	5.7	1.3
ChatGPT helped me learn the topic.	5.4	1.0
I preferred to ask ChatGPT questions instead of my professor.	5.3	1.3
ChatGPT was a good teaching assistant.	6.0	0.8

Table 2: Helpfulness of ChatGPT on the Self-directed Learning Assignment, 1=Strongly Disagree, 7=Strongly Agree, n=7

Students assessed ChatGPT on several dimensions related to technology adoption. The results are shown in Table 3. ChatGPT scores very high on ease of use and usefulness—key drivers of adoption. ChatGPT also achieved a high friendliness score. Students agreed that ChatGPT gave correct information, but a mean of 5.4 on a 7-point scale indicates that there is some room for increased accuracy. The results indicate that most students encountered some technical problems at some point.

Students answered questions about their intentions to use ChatGPT to learn topics after graduating. The results in Table 4 indicate that students will continue to use ChatGPT to select topics to learn, find learning resources, and learn the topics. All responses were “Agree” to “Strongly” agree, indicating that every participant plans to use ChatGPT for self-directed learning in the future.

Prompt	Mean	SD
ChatGPT was useful.	6.6	0.5
ChatGPT was easy to use.	6.7	0.5
I encountered technical problems while using ChatGPT.	4.9	1.6
ChatGPT was friendly.	6.6	0.5
ChatGPT gave me correct information.	5.4	1.0

Table 3: Opinions on ChatGPT Generally, 1=Strongly Disagree, 7=Strongly Agree, n=7

Prompt	Used ChatGPT (n=7)		Did Not Use (n=2)	
	Mean	SD	Mean	SD
I will use ChatGPT after I graduate to help me choose learning topics to pursue.	6.1	0.7	5.5	0.7
I will use ChatGPT after I graduate to help me understand topics I need to learn	6.0	0.8	6.0	1.4
I will use ChatGPT after I graduate to help me find learning resources.	6.3	0.8	5.5	0.7

Table 4: Anticipated Future Use of ChatGPT to Support Self-directed Learning, 1=Strongly Disagree, 7=Strongly Agree

Qualitative Results

The survey asked students to reflect on their experience and provide written feedback. First, students described their experience using ChatGPT to select a learning topic. One student reported not using ChatGPT at all. Another **student said**, “Chat GPT has helped me come up with ideas and find books for topics that are slightly specialized in a way that otherwise wouldn’t be easy to find.” **Another reported**, “It’s really good at helping you get the ball rolling with topics to pick from.”

Students reflected on ChatGPT’s usefulness for finding learning resources. Results were mixed. **One student said**, “[ChatGPT gave] some good ideas, but mostly really general obvious answers like finding a YouTube video.” **Another said**, “It does a really good job at linking viable resources to learn from and can tell you what order you

might want to learn them in. I feel like this is the main reason I will use ChatGPT in the future." **A third responded**, "It was ok, I mostly found my resources myself as well." **A student described the process by saying**, "I ask for suggestions, then for comparisons on the resources I might like, and then I ask it to clarify on the points that I care about."

The survey then asked students to reflect on **ChatGPT's support in carrying out the learning plan**. The responses showed that students critically evaluated their interactions with **ChatGPT**. **One student said**, "When I have a specific question that I can't find elsewhere, I ask ChatGPT. Even when it's wrong some of the words it uses can be crafted into a better Google search." **Another stated**, "[ChatGPT is a] useful aid for idea generation and helping troubleshoot problems, but not necessarily going to teach you." **A third said**, "It has done alright, I thought I would be able to follow the learning plan it provided step by step but I find that a lot of stuff is missing most of the time and it hasn't been consistent. At least this has been the case with anything besides code/scripts."

When asked to provide any additional feedback on using ChatGPT on the assignment, students **said**, "ChatGPT improved my learning," "Frequently unavailable," and "It's pretty great, at least so far."

6. DISCUSSION

The quantitative and qualitative results paint a fascinating picture of AI tool use for self-directed learning. Overall, the results suggest that AI tools can enhance self-directed learning, but students must be aware of potential inaccuracies provided by AI tools.

Students embraced ChatGPT but did not rely heavily on it. They forgave ChatGPT for occasionally giving incorrect and inconsistent information. One would suppose that information accuracy would be a top concern for a learning aid, but usefulness was rated higher (6.6/7) than information accuracy (5.4/7).

From an academic integrity perspective, no student engaged in any unethical behavior. Several factors could explain this result. First, students chose their own topics and so the internal drive to learn something that interested them was likely higher than if topics had been assigned to them. Second, ChatGPT could not easily create the learning evidence provided in the reflections. Students included screenshots of

applications they build, custom study guides, screenshots of course completions, and more. Third, students presented their work verbally in a peer-to-peer learning environment—something that ChatGPT could only prepare students for by helping them learn.

This self-directed learning assignment was given in a capstone course. Students had completed several previous courses in the domain and therefore had a solid foundation of knowledge. These students were better positioned to identify inaccuracies and critically interpret ChatGPT responses. Novices may have struggled more to **determine when they should challenge ChatGPT's responses**.

Two limitations of the current study should be acknowledged. First, the sample size is small. The study should be replicated to validate the consistency of the findings. Second, no objective **measurements of students' abilities to detect incorrect information from ChatGPT** were made. Future research should focus on this question could compare abilities of novice and advanced students to determine the correlation between skill level and ability to detect incorrect information.

7. CONCLUSION

The rapid pace of AI tool development and adoption creates challenges and opportunities for educators. Generative AI produces novel work that makes academic integrity enforcement challenging. Educators must adapt to the changing technological landscape not only to guard against cheating, but to improve learning outcomes.

The current study sheds light on how AI tools can be used to enhance learning. In a self-directed learning assignment, students were encouraged to use ChatGPT to aid them in developing a learning proposal for a topic of their interest, carrying out the learning plan, and reflecting on the learning process. While students noted occasional technical issues and instances of incorrect information, they reported that ChatGPT was easy to use and useful. Every student in this study plans to use ChatGPT for self-directed learning in the future.

Several aspects of the course assignment used in this study support academic integrity and rigor. First, students should be less motivated to cheat because they have the autonomy to choose their own learning topics. The ability to select learning topics should foster a genuine interest in learning.

Second, the assignment is composed of several linked milestones—a learning plan, the learning phase, a presentation, and a reflection. If a **student adopts an AI's suggestion for learning objectives**, the student is now required to meet those objectives in **the next phase. "Cheating" by developing complex learning objectives merely increases the work required in the second phase of the assignment.** Likewise, relying too much on AI in the learning phase would leave the students unprepared to present effectively or provide evidence of learning in the reflection submission. Third, the peer-to-peer presentations serve as a kind of self-administered oral examination.

Instructors should prepare students for using AI tools by creating clear assignment guidelines and encouraging students to critically evaluate AI tool output. As AI tools become integrated in more platforms, it is important that students be able to harness their power while avoiding potential pitfalls. It is quite likely that the most productive professionals will be those who judiciously embrace AI to solve problems, so AI tools should be integrated into curriculum.

8. REFERENCES

- Blöse, A. (2023, April 12). As ChatGPT Enters the Classroom, Teachers Weigh Pros and Cons. *neaToday*. <https://www.nea.org/advocating-for-change/new-from-nea/chatgpt-enters-classroom-teachers-weigh-pros-and-cons>
- Brittain, B. (2023, February 6). Getty Images lawsuit says Stability AI misused photos to train AI. *Reuters*. <https://www.reuters.com/legal/getty-images-lawsuit-says-stability-ai-misused-photos-train-ai-2023-02-06/>
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., & others. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33, 1877–1901.
- Campbell, M., Hoane Jr, A. J., & Hsu, F. (2002). Deep blue. *Artificial Intelligence*, 134(1–2), 57–83.
- Chayka, K. (2023, February 10). Is A.I. Art Stealing from Artists? *The New Yorker*. <https://www.newyorker.com/culture/infinite-scroll/is-ai-art-stealing-from-artists>
- Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2023). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, 1–12. <https://doi.org/10.1080/14703297.2023.2190148>
- Ford, B. (2023, May 1). IBM to Pause Hiring for Jobs That AI Could Do. *Bloomberg*. <https://www.bloomberg.com/news/articles/2023-05-01/ibm-to-pause-hiring-for-back-office-jobs-that-ai-could-kill>
- Fried, I. (2023, April 7). Sal Khan explains why GPT-4 is ready to be a tutor. *Axios*. <https://www.axios.com/2023/04/07/sal-khan-chatgpt-gpt4-tutor>
- Goodhue, D. L., & Thompson, R. L. (1995). Task-Technology Fit and Individual Performance. *MIS Quarterly*, 19(2), 213–236. <https://doi.org/10.2307/249689>
- Haenlein, M., & Kaplan, A. (2019). A Brief History of Artificial Intelligence: On the Past, Present, and Future of Artificial Intelligence. *California Management Review*, 61(4), 5–14. <https://doi.org/10.1177/0008125619864925>
- Hu, C. (2023, February 2). ChatGPT sets record for fastest-growing user base—Analyst note. *Reuters*. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>
- Israelsen, A. (2023, March 22). How to use ChatGPT to write code. *Learning and Development*. <https://www.pluralsight.com/blog/software-development/how-use-chatgpt-programming-coding>
- Katz, D. M., Bommarito, M. J., Gao, S., & Arredondo, P. (2023). Gpt-4 passes the bar exam. *Available at SSRN 4389233*.
- Kelly, S. M. (2023, January 28). *Real estate agents say they can't imagine working without ChatGPT now*. CNN Business. <https://www.cnn.com/2023/01/28/tech/chatgpt-real-estate/index.html>
- Klein, A. (2023, May 23). New York City Does About-Face on ChatGPT in Schools. *Education Week*. <https://www.edweek.org/technology/new-york-city-does-about-face-on-chatgpt-in-schools/2023/05>
- Kopelman, M. D. (1987). Two types of confabulation. *Journal of Neurology, Neurosurgery & Psychiatry*, 50(11), 1482–1487. <https://doi.org/10.1136/jnnp.50.11.1482>
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Longoni, C., Fradkin, A., Cian, L., & Pennycook, G. (2022). News from generative artificial intelligence is believed less. *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 97–106.
- Malik, A., Khan, M. L., & Hussain, K. (2023). How is ChatGPT Transforming Academia? Examining its Impact on Teaching, Research,

- Assessment, and Learning. *SSRN Electronic Journal*.
<https://doi.org/10.2139/ssrn.4413516>
- Marcus, G. (2022, December 19). AI Platforms like ChatGPT Are Easy to Use but Also Potentially Dangerous. *Scientific American*.
- Marquardson, J. (2020). Encouraging Lifelong Learning through Tech Explorations. *Information Systems Education Journal*, 18(6), 28–37.
- McCarthy, J., Minsky, M., & Shannon, C. (1955). *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*.
- Reddington, C. (2023, May 9). How companies are boosting productivity with generative AI. *GitHub Blog*. <https://github.blog/2023-05-09-how-companies-are-boosting-productivity-with-generative-ai/>
- Roose, K. (2023, January 12). Don't Ban ChatGPT in Schools. Teach With It. *New York Times*.** <https://www.nytimes.com/2023/01/12/technology/chatgpt-schools-teachers.html>
- Rosenzweig-Ziff, D. (2023, January 5). New York City blocks use of the ChatGPT bot in its schools. *The Washington Post*. <https://www.washingtonpost.com/education/2023/01/05/nyc-schools-ban-chatgpt/>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, \Lukasz, & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
- Venkatesh, V., Thong, J. Y. L., & Xu, X. (2012). Consumer Acceptance and Use of Information Technology: Extending the Unified Theory of Acceptance and Use of Technology. *MIS Quarterly*, 36(1), 157–178.
- Verma, P. (2023, May 18). A professor accused his class of using ChatGPT, putting diplomas in jeopardy. *The Washington Post*. <https://www.washingtonpost.com/technology/2023/05/18/texas-professor-threatened-fail-class-chatgpt-cheating/>
- Weiser, B. (2023, May 27). Here's What Happens When Your Lawyer Uses ChatGPT. *New York Times*.** <https://www.nytimes.com/2023/05/27/nyregion/avianca-airline-lawsuit-chatgpt.html>
- Zhang, Y., Li, Y., Cui, L., Cai, D., Liu, L., Fu, T., Huang, X., Zhao, E., Zhang, Y., Chen, Y., Wang, L., Luu, A. T., Bi, W., Shi, F., & Shi, S. (2023). *Siren's Song in the AI Ocean: A Survey on Hallucination in Large Language Models*. <https://doi.org/10.48550/ARXIV.2309.01219>

APPENDIX A Sample ChatGPT Interaction

A student supplied the following ChatGPT interaction.

Chat GPT Prompt

```
what does this TCP stream mean: 220 x.shh.sh ESMTP Exim 4.86 Wed, 06 Sep 2017 22:11:43 +0000
ehlo x.shh.sh 250-x.shh.sh Hello x.shh.sh [::1] 250-SIZE 52428800 250-8BITMIME 250-PIPELINING
250-AUTH LOGIN 250-STARTTLS 250-PRDR 250 HELP auth login 334 VXNlcm5hbWU6bWFSbG9yeQ==
334 UGFzc3dvcmQ6ZmxhZy1zcGluc3Rlci1iZW5lZmI0LWZhbHNpZnktZ2FtYmlhbg== 535 Incorrect
authentication data 421 x.shh.sh lost input connection
```

Chat GPT Response

This TCP stream appears to be a sequence of commands and responses exchanged between an email client and an email server using the SMTP (Simple Mail Transfer Protocol) protocol. Here's a breakdown of what's happening:

1. The email client sends the command "ehlo x.shh.sh" to the server. This command is used to initiate a session with the server and to identify the client to the server.
2. The server responds with "250-x.shh.sh Hello x.shh.sh [::1]", indicating that it has received the client's greeting and is ready to proceed with the session. The server also provides a list of capabilities it supports, including "SIZE 52428800", "8BITMIME", "PIPELINING", "AUTH LOGIN", "STARTTLS", "PRDR", and "HELP".
3. The email client sends the command "auth login" to the server, indicating that it wants to authenticate itself to the server using the LOGIN method.
4. The server responds with "334 VXNlcm5hbWU6", which is a Base64-encoded message that prompts the client for its username.
5. The email client responds with "bWFSbG9yeQ==", which is a Base64-encoded version of the client's username "malory".
6. The server responds with "334 UGFzc3dvcmQ6", which is a Base64-encoded message that prompts the client for its password.
7. The email client responds with "ZmxhZy1zcGluc3Rlci1iZW5lZmI0LWZhbHNpZnktZ2FtYmlhbg==", which is a Base64-encoded version of the client's password "flag-spinner-benefit-falsify-gambian".
8. The server responds with "535 Incorrect authentication data", indicating that the client's login attempt failed.
9. The server sends the message "421 x.shh.sh lost input connection", indicating that the connection has been lost and the session is over.

Based on this TCP stream, we can conclude that the email client attempted to authenticate itself to the server using the LOGIN method, but the authentication failed because the provided credentials were incorrect. We can also see that the server supports a number of other capabilities, such as message size limits, MIME encoding, and encryption.