

Enhancing Professional Military Education with AI

Best Practices for Effective Implementation

Adam T. Biggs

U.S. Army Command and General Staff College

Abstract

The increasing prevalence of artificial intelligence (AI) tools has significant implications for professional military education (PME). As AI technologies continue to evolve, they offer new opportunities to enhance learning outcomes, improve educational efficiency, and support the development of critical skills for military professionals. However, integrating AI tools in PME also raises important questions about their effective use, ethical considerations, and potential pitfalls. This article provides an overview of the current state of AI tools in PME, discusses the benefits and challenges associated with their use, and offers best practices for military educators to optimize their implementation. By examining the potential applications and limitations of AI tools, this article will inform the development of effective strategies for leveraging AI in PME, ultimately enhancing the learning experience and preparing military professionals for the complexities of the modern operational environment.

Clear communication is a critical skill enhanced during professional military education (PME). As with any critical skill involving a constant pursuit of improvement, many different tools and techniques exist to foster better development of professional communication. Artificial intelligence (AI) is one such option that has garnered significant attention for its potential and its boundless innovation. These AI products have become nearly unavoidable and touch every facet of daily life. For example, educators can now use AI-powered platforms to streamline different administrative functions to include grading essays (Chen et al., 2020), and healthcare professionals are seeking new ways to utilize AI for clinical decision-making (Secinaro et al., 2021). Big tech companies like Google and Microsoft have further invested billions of dollars to integrate AI tools into their existing product lines

(Rattner, 2024). With so much investment into AI products, their seeming omnipresence will likely become a permanent reality.

PME could likewise benefit from further integration of AI into the curriculum. Impending changes are evident as accredited and degree-granting programs within the PME space have altered policies to permit the use of AI for coursework (U.S. Army Command and General Staff College, 2024). In practice, students could benefit from these AI tools in many ways. Someone might upload a draft essay and ask for feedback on their current version, along with recommendations for improvement. Artwork could be generated to support classroom war-gaming or key presentations. Students might also brainstorm writing ideas or generate summaries of articles through AI. Whatever the specific problem set, there is likely some application where a student could integrate an AI tool into their military studies.

Despite lifting prohibitions against AI for PME, there remains a great deal of exploration that must occur. For example, even if personnel can use AI, what are the optimal applications? How do service members implement AI to enhance their professional communication without violating ethical standards? What military-specific challenges might arise in a military context that would not apply as readily to other areas of education? After all, military jargon alone might provide a hurdle to a large language model (LLM) if the base text builds upon civilian dialogue—and this concern exists alongside the obvious security issues of entering military orders, data, and documents into sometimes publicly available AI products that learn from the information entered. There are many concerns as to how students should use AI products in PME, yet there is also ample opportunity to begin developing best practices to support effective and ethical AI use during coursework. Higher echelons have already issued some guidance regarding military applications for AI products (U.S. Training and Doctrine Command, 2024).

Although AI tools could greatly enhance PME, effective implementation requires understanding both their capabilities and limitations. The current discussion will identify some best practices and ethical pitfalls when integrating AI into professional military studies. As such, the key goal is to enhance communication among future military leaders while educating them on the challenges of AI tools. The discussion will begin by describing the development of AI programs and LLMs that have received recent popularization as new commercial products. The first objective is to establish a base understanding for individuals about how AI products are developed

Lt. Cmdr. Adam T. Biggs is a research psychologist in the Medical Service Corps of the U.S. Navy. His primary duties involve conducting translational research to support fleet needs in a variety of human performance, organizational, and medical research areas. In the completion of these duties, Biggs has published numerous research studies while also training in various research ethics topics. He holds a BS from McKendree University, and an MA and PhD from the University of Notre Dame.

and their general capabilities. Next, the focus will shift to ethical pitfalls and key problems that exist when utilizing AI tools. Finally, the discussion will review best practices on how AI can be used to support PME. The goal for reading should be to supplement comprehension when dealing with larger-than-normal reading volume, and the goal for writing should be to enhance professional communication while appropriately crediting sources and AI material without violating academic policies. Ultimately, the collective discussion aims to enhance PME by building upon recent changes in policy that allow AI products to enhance the educational experience.

Artificial Intelligence: What Is “AI” and How Can It Be Used for Military Applications?

In recent years, AI evolved from an abstract topic of cognitive science (cf. Fetzer, 1990) to a dominant force among commercial, educational, and industrial sectors. A multitude of products now incorporate AI or claim AI development to tout their enhanced potential. Businesses have integrated AI into their product lines to deliver better solutions for customers, and noncommercial entities have likewise sought to utilize this technological enhancement in their respective spheres of influence. For example, people have explored AI integration for diverse applications such as natural disaster responses (Sun et al., 2020) and medicine (Meskó & Görög, 2020). Still, despite the seeming omnipresence of AI solutions in daily life, this technology remains under continuous development with many people retaining only a cursory understanding of it. Therefore, the first question must remain the obvious one: What is AI?

Many people use the term “AI” as a catchall for metaphors, mental models, and word prediction paradigms without a common definition for what does or does not qualify as AI (Heaven, 2024). Contrary to popular usage, most current AI models exploit LLMs rather than true AI. An LLM processes enormous volumes of data to learn patterns and adjust feedback to approximate a human response (Zhou et al., 2024). Essentially, an LLM predicts what humans would say by examining large volumes of text to identify predictable patterns. There is no true intelligence to the response, merely a probable combination of outputs. That said, the models can become increasingly more reliable with larger and larger input, which previously limited their dependability as a function of computer processing power through both initial training and data available to process.

Technological advances have solved a substantial portion of the problem as smaller and smaller computers have larger and larger computing power. In this way, LLMs have become capable of processing enough predictable relationships to approximate realistic human responses, hence the oft-mislabeled distinction as AI when the real description should be LLMs. True AI is instead known as artificial general intelligence (AGI; McLean et al., 2023). The key distinction is the capacity to transfer

learned knowledge and processes to new domains rather than being restricted only to the learned domain. Although AGI remains a theoretical concept, for the time being, this capability could adapt to new environments without the explicit programming to support novel applications.

Despite remaining limited to machine learning and LLMs, many systems called AI have developed remarkable capabilities when responding to user inputs. Search engines now regularly include AI overviews as summaries for certain queries. Likewise, reading platforms frequently come accompanied by AI tools to aid in summarizing or processing the main text. These tools have also reached a point of maturity where the outputs cannot be ignored as simplistic or trivial. Modern AI platforms continue to refine output with increasingly more meaningful capabilities. In turn, there is potential to utilize these tools for support in higher education, and a growing number of advocates argue for permitting AI tools in PME (Kelly & Smith, 2024).

When applied to the military context, there are a few important considerations to note that make AI usage for PME or military-specific AI tools different from other forms of technology. Foremost, an LLM predicts text based upon relationships learned from a preliminary training stage. ChatGPT, a large natural language processing algorithm, incorporated 570 gigabytes of data in its training phase (Heikkilä, 2023). Even a conservative evaluation would suggest this volume of data includes hundreds of thousands or millions of texts and billions of words. Nevertheless, the learning dataset is also a restriction unto itself. AI models depend upon the text used during their training phase to make predictions about the next word or when evaluating content. For military applications, the training set becomes a double-edged sword. Any generalized training data might not be capable of addressing military parlance or problems, and there would be massive operational security violations to train a widely available resource with military data. Specifically, if a publicly available LLM were supplemented with military data for further training, anyone with access could ask questions that reveal information from data reviewed during the training phase. Adversaries could peruse controlled military documents at will through this vulnerability.

Instead, the solution is to develop controlled military datasets for training military-specific AI tools. These instruments can be constrained to specific information that best exemplifies the military context by uploading only military sources. Such tools would need to be restricted and limited to the unclassified or classified systems on which they learned. Even so, this limitation is no more restrictive than any other constraint accompanying classification for operational purposes. More importantly, the Department of Defense has already begun building and deploying AI tools for military purposes, and the reception has been voracious. The U.S. Air Force and Space Force released an AI tool for internal use dubbed NIPRGPT (the Non-Classified Internet Protocol Generative Pretraining Transformer) in 2024; three months after its release, over 80,000 airmen and guardians experimented with the system

(Albon, 2024). Perhaps the most important lesson from this context is the inevitability of AI tools. Service members will encounter them in daily life, and they will be eager to employ these tools in their professional duties.

Ethics and Challenges in Using AI Tools

The most straightforward ethical issue comes from a simple assumption—namely, that the output of AI tools is precisely what it purports to be. Too many people presume that the answer to a prompt is factual. However, AI can “hallucinate,” which describes how AI might generate highly skewed, misleading, or outright false content (Lakhani, n.d.). There is no singular reason why hallucinations occur. Some instances might be due to biased training data, outdated information, or a model attempting to overfit a response based on what it has learned. The latter example can produce even purely fictitious claims if the model training involved recognizing and processing certain formats. Still, an important thing to consider is that AI tools are designed to provide a response. Whereas a student might admit not knowing an answer, the AI tool will provide something whether that response represents accurate information or not. Viewed in this light, hallucinations are a byproduct of an algorithm programmed to provide a response whenever prompted. The inherent danger is assuming the output to be factual.

Among the various instances of hallucinations catching people off-guard, there is an example of how damaging the assumption of accuracy can be. In a 2023 New York aviation lawsuit, attorneys utilized ChatGPT to help them prepare a federal court filing, which they presented to the court as the AI tool had delivered (Bohannon, 2023). Unfortunately, the program hallucinated and produced not one, but six fictitious cases to show precedence for their claims in court. When discovered, the judge eventually sanctioned the attorneys for dereliction of their responsibilities by presuming the cases were real and not investigating the cited precedence themselves (Merken, 2023). Moreover, they are no longer alone in this embarrassment. Other cases have occurred where lawyers have allegedly used AI tools to prepare cases without properly investigating the outcome, only for the AI to hallucinate and cite more nonexistent cases (Cecco, 2024). These examples represent actual cases where individuals who accepted information without verifying the record faced severe real-world consequences.

Another challenge involves AI translations between languages. Neural machine translation, among other techniques, has greatly enhanced the accuracy of translations through supporting software (Mohamed et al., 2024). AI tools have been remarkable in advancing this capability. However, the translations are not perfect and misunderstandings can cause severe consequences. For example, people have been denied asylum in some cases because translation errors misrepresented their case to

immigration authorities (Bhuiyan, 2023). This instance would appear to be an ideal use case, where limited border authorities could utilize technology to cover shortages in manpower while still addressing the many language-related issues that could arise in border crossings. Instead, the example demonstrates how subtle differences in meaning that an interpreter might catch could be overlooked by AI software. Nuance becomes one of several possible underlying explanations for the discrepancy. Specifically, language-learning systems apply better to high-resource languages with many examples for input, like English or Chinese, but might encounter significant problems converting from English to other languages (Gordon, 2024). AI tools are becoming more robust each day, yet they currently lack the capability to parse nuance the way a human might.

The learning set itself could be a problem that leads to ethical misunderstandings. In an academic environment, plagiarism is a common concern wherein one student takes credit for someone else's work. Previously, plagiarism would become an issue when students copied from someone else or failed to cite appropriately throughout their writing, but AI tools introduced a new wrinkle to this problem. Because AI tools often learn from prompts and material with which they interact, the same algorithms could learn from related work and provide answers that seem original without being so. Students may believe the work to be an original AI generation, and therefore they would not be plagiarizing an individual. Nonetheless, AI may be regurgitating related work from which it learned that too closely approximates text from another student. This possibility is a problem for any black-box-style learning system, which describes a system or process where the inner workings lack transparency. Black box learning instead relies upon input-output relationships, whereas internal learning procedures cannot be fully documented or generated. Simply put, no one may fully understand why an AI product generated a given response because they cannot fully replicate the logic developed during its training.

Even if proper citation could address the plagiarism problem, citing AI usage differs from a typical citation. Other media or scholarly sources have some method to identify the author or organization when citing the originating idea; yet AI tools generate the information without an independent author to cite. This issue too has led some students to believe that work generated with AI does not require citations. To avoid the issue entirely, many universities have adopted new methods for properly citing and crediting AI tools when used to develop research or other written products (Brown University Library, 2025). The intent is merely to ensure that instructors can appropriately gauge critical thinking in writing, or in the case of research efforts, the authors provide a reproducible pathway to identify sources.

Some PME programs have likewise instituted policies in accordance with these ideas that permit the assistance of AI tools in writing (e.g., U.S. Army Command and General Staff College, 2024). That said, students must continue to submit original work for educational assessments, which is why there must be some understand-

ing as to what the student generated without help and what elements could be attributed to AI assistance. Student guidance currently identifies that AI tools could be helpful in analyzing writing prompts, assembling outlines for class writing assignments, summarizing source material, and offering suggestions in editing (Lythgoe et al., 2024). Any one of these options represent powerful tools to help writers produce higher-quality material, especially if they have not produced scholarly work in some time. The caveat is merely to ensure that students cite all AI-generated content through footnotes that document prompts or other edits as contributed by AI tools (Lythgoe et al., 2024).

Furthermore, footnote entries offer an interesting middle ground to the challenge of original work assisted by AI for multiple reasons. Foremost, there is a process to identify how a student utilized AI, which is important since there is technically no original source to cite. This documentation limits the extent of confusion that might result if academic integrity checks flag material as unoriginal or plagiarized content. Additionally, footnotes are not intended to be lengthy accounts within a manuscript the way endnotes or appendices might be. A footnote provides an opportunity for documentation while inherently limiting the space available. As a general rule, if capturing AI support becomes cumbersome enough to warrant a full appendix, the individual is probably relying too much upon AI for content generation.

Thus far, these ethical issues have largely resulted from accepting AI output at face value or falsely claiming content generated through AI as original work. Other issues that arise within a research context concern the unintentional infringement of individual rights. Specifically, research ethics provides many different tools to protect the rights of research participants. These rules include ethical oversight and informed consent if the research involves human subjects. When involving AI tools, there is the potential for private information to be released or for available information to become identifying when presented in aggregate, thereby raising privacy concerns when using generative AI (University of North Carolina, n.d.). For example, someone might enter research data into a publicly available AI tool that learns from updated information. The details might be de-identified when contained, but the uploader cannot know everything else processed through the platform. If the system encounters related information, there is a possibility of integrating old data and new data into a learning model that produces spillage. In essence, entering data (including datasets, unpublished work, or other proposals) into public places is tantamount to public release, and the uploader cannot predict how the AI tool will process or distribute this information. This unknown creates a potential vulnerability for individual privacy. Universities, publishers, and funding organizations are trying to catch up with the emerging AI tools for research applications, and in the short-term, there are significant ethical considerations for AI in research around which these organizations are still developing norms, requirements, and best practices.

Another area of concern becomes apparent in the blended use of AI with two different users applying AI for complementary functions. Basically, if one user generates content with AI, and another user processes that content with AI, there is the potential for a portion of their interaction to become dominated by AI processing rather than human interaction. One user could knowingly attempt to gain an advantage through manipulating content if they know the user on the other end will utilize AI to process their content. A prime example of this possibility is a human resource manager using AI tools to help sift through many résumés for a particular job opening. There are tricks people have employed to gain an advantage in this context, such as blending white font into the document (Abril, 2023). Human readers would not see the white text without further inspection, yet AI would process it the same as any other text. Someone could use the opportunity to insert numerous keywords aligned with the job opening to raise interest in their application. Alternatively, someone could enter commands for the AI to secure a desired outcome. In the case of AI-assisted interviewing, the applicant could instruct the AI tool to tell hiring managers that they are the ideal fit for a job. Whether these actions are truly unethical or a novel business practice to garner attention, in an academic environment, the concern is students using AI to circumvent their instructors. This situation could arise if instructors are using AI tools to assist in their evaluations of student work. As such, the example is an important demonstration that instructors should be careful when using AI to avoid unintended consequences.

Tips and Tricks to Effectively—and Ethically—Use AI Tools in Professional Military Education

LLMs and AI have the potential to enhance education in numerous ways, including through the production of novel educational content, to enhance student engagement, and to personalize learning experiences (Kasneci et al., 2023). Understanding ethical challenges helps lay the groundwork for effective usage of AI tools in PME, though this information does little by itself to optimize AI use. Instead, there are several tips and tricks developed by ambitious people over the past few years that could help students maximize their possible benefits from these technological instruments. There is an inherent focus on AI support of writing in the following advice. Nevertheless, there is overlap in applying these tips for AI as a study tool as well (see Table 1 for an overview).

Before considering more advanced use of AI tools, the first tip applies to beginners. Summaries and background information are two things AI normally processes well since the task merely involves presenting facts. However, AI tools—at present—cannot produce human-level understanding and synthesis of information. Thus, the first beginner mistake is to ask simple questions of AI tools and develop only a

Table 1*AI Challenges and Best Practices*

Key Problem or Ethical Challenge	Best Practice to Implement
AI can hallucinate and provide false or misleading answers.	Ask AI to provide sources for its information. Independently verify sources rather than take AI output at face value.
Accusations of plagiarism if AI provides answers similar to previous data.	Document prompts and interaction with AI programs to maximize replication potential and minimize false claims of plagiarized work.
Students rely solely upon AI rather than reading the required texts.	Use AI to supplement reading rather than as a replacement. Complete a read through first to provide a base understanding before asking AI questions.
Translations and nuance not always captured by AI software.	Ask multiple iterations of the same question to capture nuance in meaning. Use different perspectives when framing questions to capture the problem set from different angles.
Entering military data or military-relevant prompts into a large language model with public access.	Treat commercially available AI tools as belonging to the public domain. Avoid creating operational security vulnerabilities by entering controlled information into public domain. Use constrained notes, restricted datasets, and military-specific AI tools when conducting official military functions.
Data entry should be considered public release of information and could jeopardize privacy.	Assume AI will aggregate data and potentially violate individual privacy for personal information. Consult the Human Research Protections Office for further guidance to avoid confusion.
People can manipulate materials if they believe AI will be used to process the information.	Ask how someone might use AI to process your materials to “red team” their perspective. Visually inspect products before submitting them to AI analyses.

superficial understanding of how AI answers questions. A more effective beginner technique is to explore AI with a topic or document the student already knows quite well. Because the student already has an ample foundation of knowledge, they will better parse the flaws, nuance, and limitations of AI by exploring the tools with these texts. For example, a student who is an avid baseball fan might ask ChatGPT questions about what a baseball player should do in given situations or ask an AI product to provide historical comparisons of different famous players. The intent is to help the student learn the rhythms and responses that AI tends to give. By beginning with source material the student knows well, AI peculiarities become more apparent. This understanding helps the student later since they will limit their use of answers received from AI accordingly.

Dealing with hallucinations represents the most obvious concern in trusting AI responses. After all, if neither students nor staff could trust the output of AI tools, what purpose could there be in seeking their support? While continued innovation should limit the possibility for outright hallucinations to corrupt results in the future, several different techniques can minimize the challenges posed by AI hallucinations today. These possibilities include asking the AI to provide sources (or otherwise have some means of identifying evidence to support the answer), entering multiple prompts to contrast the output, inquiring AI for its reasoning behind an answer, and most importantly, double-checking the output information independently (Lakhani, n.d.). The last point is the most important, and it works best when asking AI to cite sources. Students can then independently verify whether the information produced is accurate. Granted, this advice is good for any potentially biased output, whether belonging to internet media or AI-generated content. Different AI platforms will readily adapt to outputting references in different preferred formats, yet the true opportunity is the chance to follow up and determine if the information appears legitimate. This step also requires less work than it might seem. Many scholarly search databases such as Google Scholar and PubMed index millions of scholarly articles. If the purported citations cannot be verified through one of these platforms, then the student should grow increasingly skeptical that the output might be an AI hallucination.

Two other possibilities address hallucinations and uncertain information through complementary approaches. First, multiple prompts allow the student to assess information reliability through consistency. This method does not mean simply rephrasing a question using different words. Instead, try approaching the question from another perspective. Some AI tools benefit from different personas that enable answering from another point of view. Consider an example of a leadership case study involving police reform in New York City (cf. Kim & Mauborgne, 2003).¹ A standard approach would involve asking an AI tool to summarize this article. Alter-

¹ This case study is also used in a U.S. Army Command and General Staff College leadership course (L100).

natively, the student could provide different contexts when asking AI for information by adopting a different perspective each time. This scenario would allow students to ask questions from the perspective of a police commissioner, an officer patrolling the streets, the mayor's office, media outlets, or even criminals. Each perspective should have different answers to certain questions, especially as attitudes and effectiveness of crime prevention techniques would be concerned. Adaptation to each question helps limit not just the possibility of hallucinations influencing the outcome, but this approach also creates a more holistic understanding of the situation. Moreover, this method enables another strategy to avoid hallucinations and dive deeper by asking AI its reasoning behind the information provided. Both law enforcement and criminal perspectives should give similar answers about basic facts such as dates within the story, yet each perspective should have different reasoning underlying its response to the intrusiveness of crime prevention techniques. If comparable answers are given for both, then the similarities should be a red flag that the student cannot fully trust the AI output, that the synthesis of information is marginal at best, or other possibilities that warrant a deeper dive into the material before accepting AI results.

Another technique is to utilize retrieval augmented generation (RAG; Rogers, 2024). RAG searches constrain the possible answers to a set of real documents to limit the possibility of hallucination. This technique could utilize a set of consolidated notes to limit the possible input or engage a search engine to pull in real documents. Granted, the AI prompt must further anchor responses only to the identified subset of documents and not all material encountered during initial model training. The latter possibility creates an opportunity for misleading results despite an active effort to avoid hallucinations because it remains reliant upon accepting AI output as genuine. Success thus depends on how effectively the AI tool can narrow focus only to relevant information without drawing upon its initial training or information outside the constrained set—essentially keeping an onus on the searcher to construct an effective prompt while narrowing the existing documents to be searched. As such, RAG does add value and limits hallucinations, although the output information would still benefit from citations, sources, or another means of confirming that the information is indeed genuine. Furthermore, there are a few different names for this technique. Some outlets might call it consolidated notes or related language describing the limited search parameters. Nonetheless, the important element is that answers become limited to a particular set of information rather than asking the algorithm to draw upon all previous facts and information it might have encountered.

Further techniques should only be employed once the student has developed some mastery with AI tools. Although these techniques unlock the greatest potential for AI assistance in military education, they also involve the most nuance and therefore require some base level of familiarization before they can be fully utilized. In short, these techniques allow prompt engineering to maximize AI outputs. Prompt

Table 2

Prompt Examples and Best Practices

Prompt	Key Points
"Tell me what I should know about this article."	<p>Vague or incomplete request.</p> <p>No guidance about the type of information requested.</p> <p>No template given to structure the output into a more usable format.</p> <p>No detail about the desired length of response.</p> <p>Wording is not precise and could lead to confusion.</p>
"Analyze the positive and negative elements of this article using the SWOT (strengths, weaknesses, opportunities, and threats) framework with a background section for broad overview and conclusion section of less than 500 words. The target audience for the summary is a classroom of students in professional military education at a graduate level."	<p>Direct request with active voice.</p> <p>Specifics given about the type of information requested.</p> <p>Example format provided to structure the response.</p> <p>Detailed requirements for length.</p> <p>Precise wording likely to deliver clear results.</p> <p>Tailored to a specific audience.</p>

engineering is the skill inherent to crafting questions that produce optimal outputs when entered into AI tools (Snow, 2023; see Table 2 for examples).

Every prompt will inherently have some sort of task or command since the user is asking an AI tool to do something. That said, not every prompt achieves the same quality of output, and prompt engineering becomes the art form that will differentiate individuals who excel when using AI tools and individuals who simply use their functions. For the task, military users should be familiar with the type of clear direction often recommended (Snow, 2023). An example might involve directing the AI to "summarize the key takeaways of the article," but this direction is only a starting point. Active voice helps, although tweaks could optimize the output. Layer requests by adding specific components desired in the output. For example, the same prompt could be improved by asking the AI tool to "summarize the leadership best practices in this article, include a bullet point summary of key takeaways, and provide a conclusion section of no more than 250 words." Specific requests written in an active voice help refine the task in ways that allow AI tools to produce a better output. Thus, optimal output can be achieved when describing the requested task with specificity.

Additional refinement can further augment the prompt, depending on the situation and tool in question. Some tools will benefit from examples that help provide

format and structure to the output. For example, rather than ask for the positives and negatives of a certain article, someone could frame the prompt as “analyze the positive and negative elements of this article using the SWOT (strengths, weaknesses, opportunities, and threats) framework with a background section for broad overview and conclusion section with key takeaways.” Examples help refine the requested task into a more constrained format by providing the AI tool with context upon which to craft a response. If asking ChatGPT to identify good restaurants in the area, as another possible use, constrain the response with details like the quality, cost, location, type of food, or other details critical to any decision. Of course, any military-specific usage will be restricted by the type of information that can be entered into the platform. Some current guidance outright restricts commercial, off-the-shelf AI programs for use in any professional purpose given the risks of unsecure data storage, potential for hallucinations, and lack of transparency (U.S. Army Combined Arms Center, 2024). Professional military use cases should restrict AI use to approved platforms such as NIPRGPT or CamoGPT.

Other best practices in prompt engineering include the voice used in crafting the prompt. Remember, LLMs learned from enormous datasets that included a wide range of information, sometimes presented in different contexts. A biased voice or passive voice could prompt the AI tool to seek matching style, and so the response could be equally biased or passive in response to the prompt. Moreover, emotion can further change the context. Chatbots can be primed with encouraging words to perform better, but under most circumstances, a moderate amount of politeness achieves better results than flattery or aggression when crafting a prompt (Ziegler, 2024). Professional tone is often the best example when entering prompts into AI tools. Finally, for tasks someone will need to do repeatedly, users can keep a prompt database of inputs that have been successful during previous iterations. Over time, these prompts can be developed and refined even further to maximize the interactions. This possibility might be especially important for military users who eventually employ AI tools to develop orders or other highly structured tasks with common elements between iterations.

Of course, most of the discussion focused on student use of AI tools. There are also important applications for instructor use of AI in PME. One possibility would be to help adapt the curriculum to new material. AI could generate supporting images or instructors could explore new material when developing lesson plans. These additions could help instructors shape the curriculum with feedback from AI tools. That said, the role of the instructor becomes subject to similar advice and best practices given for the student. AI can provide ideas, yet the same hallucinations and false leads could deceive instructors the same way it might have students. Instructors should likewise proceed with caution if considering AI to facilitate their grading requirements. The best practice for either curriculum development or classroom instruction would be to brainstorm with AI support while double-checking all sources for accuracy.

Summary

AI tools have greatly evolved in recent years. The concept has advanced from a novelty to a practical toolset available throughout multiple facets of daily life from supporting education to making dinner plans in new cities. For PME, there are many possibilities that students could use to further their learning as AI tools can assist with large reading requirements and writing exercises. Nevertheless, especially in a military context, there are some evident downsides. AI bots could produce misleading results when they hallucinate, or improper citation could lead to confusion and accusations of plagiarism. As much as these tools have advanced recently, their integration into educational environments remains preliminary at best. Both teachers and students are attempting to identify the best practices of using AI to support a learning environment. For those individuals who choose to utilize AI tools in PME, perhaps the three best pieces of advice right now are

1. Never accept the full output of AI tools without double-checking sources.
2. Always properly cite uses of AI in academic work.
3. AI tools are best utilized as a supplement to enhance reading and writing exercises, not as a replacement for doing the work. ✎

Disclaimer

The views expressed in this article are those of the author and do not necessarily reflect the official policy or position of the U.S. Army Command and General Staff College, Department of the Navy, Department of Defense, or the U.S. government. The author is a military service member. This work was prepared as part of his official duties. The author has no financial or nonfinancial competing interests in this manuscript.

References

- Abril, D. (2023, July 24). Job applicants are battling AI résumé filters with a hack. *Washington Post*. <https://www.washingtonpost.com/technology/2023/07/24/white-font-resume-tip-keywords/>
- Albon, C. (2024, September 16). Air Force's ChatGPT-like AI pilot draws 80K users in initial months. *Defense News*. <https://www.defensenews.com/air/2024/09/16/air-forces-chatgpt-like-ai-pilot-draws-80k-users-in-initial-months/>
- Bhuiyan, J. (2023, September 7). Lost in AI translation: Growing reliance on language apps jeopardizes some asylum applications. *The Guardian*. <https://www.theguardian.com/us-news/2023/sep/07/asylum-seekers-ai-translation-apps>
- Bohannon, M. (2023, June 8). Lawyer used ChatGPT in court—and cited fake cases. A judge is considering sanctions. *Forbes*. <https://www.forbes.com/sites/mollybohannon/2023/06/08/lawyer-used-chatgpt-in-court-and-cited-fake-cases-a-judge-is-considering-sanctions/>

- Brown University Library. (2025, February 5). *Generative artificial intelligence*. <https://libguides.brown.edu/AI/LearningCommunity>
- Cecco, L. (2024, February 29). Canada lawyer under fire for submitting fake cases created by AI chatbot. *The Guardian*. <https://www.theguardian.com/world/2024/feb/29/canada-lawyer-chatgpt-fake-cases-ai>
- Chen, L., Chen, P., & Lin, Z. (2020). Artificial intelligence in education: A review. *IEEE Access*, 8, 75264–75278. <https://www.doi.org/10.1109/ACCESS.2020.2988510>
- Fetzer, J. H. (1990). What is artificial intelligence? In *Artificial intelligence: Its scope and limits* (pp. 3–27). Springer Netherlands.
- Gordon, S. F. (2024, February 5). Artificial intelligence and language translation in scientific publishing. *Science Editor*, 47(1), 8–9. <https://www.csescienceeditor.org/article/artificial-intelligence-and-language-translation-in-scientific-publishing/>
- Heaven, W. D. (2024, July 10). What is AI? MIT *Technology Review*. <https://www.technologyreview.com/2024/07/10/1094475/what-is-artificial-intelligence-ai-definitive-guide/>
- Heikkilä, M. (2023, April 19). OpenAI's hunger for data is coming back to bite it. MIT *Technology Review*. <https://www.technologyreview.com/2023/04/19/1071789/openais-hunger-for-data-is-coming-back-to-bite-it/>
- Kasneci, E., Sessler, K., Küchemann, S., Bannert, M., Dementieva, D., Fischer, F., Gasser, U., Groh, G., Günnemann, S., Hüllermeier, E., Krusche, S., Kutyniok, G., Michaeli, T., Nerdel, C., Pfeffer, J., Poquet, O., Sailer, M., Schmidt, A., Seidel, T., ... Kasneci, G. (2023). ChatGPT for good? On opportunities and challenges of large language models for education. *Learning and Individual Differences*, 103, 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Kelly, P., & Smith, H. (2024, May 23). How to think about integrating generative AI in professional military education. *Military Review*. <https://www.armyupress.army.mil/journals/military-review/online-exclusive/2024-ole/integrating-generative-ai/>
- Kim, W. C., & Mauborgne, R. (2003, April). Tipping point leadership. *Harvard Business Review*, 81(4), 60–69. <https://hbr.org/2003/04/tipping-point-leadership>
- Lakhani, K. (n.d.) *How can we counteract generative AI's hallucinations?* Digital Data Design Institute at Harvard. <https://d3.harvard.edu/how-can-we-counteract-generative-ais-hallucinations>
- Lythgoe, T. J., Boyce, A. S., Crowson, T. A., Kalic, S. N., McConnell, R. A., Noll, M. L., & Reider, B. J. (2024, March). *Professional writing: The Command and General Staff College writing guide* (Student Text 22-2). U.S. Army Command and General Staff College. https://armyuniversity.edu/cgsc/cgss/DCL/files/ST_22-2_US_Army_CGSC_Writing_Guide_March_2024.pdf
- McLean, S., Read, G. J., Thompson, J., Baber, C., Stanton, N. A., & Salmon, P. M. (2023). The risks associated with artificial general intelligence: A systematic review. *Journal of Experimental & Theoretical Artificial Intelligence*, 35(5), 649–663. <https://doi.org/10.1080/0952813X.2021.1964003>
- Merken, S. (2023, June 26). New York lawyers sanctioned for using fake ChatGPT cases in legal brief. Reuters. <https://www.reuters.com/legal/new-york-lawyers-sanctioned-using-fake-chatgpt-cases-legal-brief-2023-06-22/>
- Meskó, B., & Görög, M. (2020). A short guide for medical professionals in the era of artificial intelligence. *npj Digital Medicine*, 3(1), 126.

- Mohamed, Y. A., Khanan, A., Bashir, M., Mohamed, A. H. H., Adiel, M. A., & Elsadig, M. A. (2024). The impact of artificial intelligence on language translation: A review. *IEEE Access*, 12, 25553–25579. <https://doi.org/10.1109/ACCESS.2024.3366802>
- Rattner, R. (2024, August 3). Breaking down the tech giants' AI spending surge. *Wall Street Journal*. <https://www.wsj.com/finance/stocks/breaking-down-the-tech-giants-ai-spending-surge-e282ca24>
- Rogers, R. (2024, June 14). Reduce AI hallucinations with this neat software trick. *Wired*. <https://www.wired.com/story/reduce-ai-hallucinations-with-rag/>
- Secinaro, S., Calandra, D., Secinaro, A., Muthurangu, V., & Biancone, P. (2021). The role of artificial intelligence in healthcare: A structured literature review. *BMC Medical Informatics and Decision Making*, 21, 1–23. <https://doi.org/10.1186/s12911-021-01488-9>
- Snow, J. (2023, April 12). ChatGPT can give great answers. But only if you know how to ask the right question. *Wall Street Journal*. <https://www.wsj.com/articles/chatgpt-ask-the-right-question-12d0f035>
- Sun, W., Bocchini, P., & Davison, B. D. (2020). Applications of artificial intelligence for disaster management. *Natural Hazards*, 103(3), 2631–2689. <https://doi.org/10.1007/s11069-020-04124-3>
- University of North Carolina. (n.d.). *Generative AI usage guidance: Research community*. <https://provost.unc.edu/generative-ai-usage-guidance-for-the-research-community/>
- U.S. Army Combined Arms Center. (2024, December 9). *Combined Arms Center guidance on generative artificial intelligence (GENAI) and large language models* (Command Policy Letter 24).
- U.S. Army Command and General Staff College. (2024, May 9). *Command and General Staff College academic ethics policy* (CGSC Bulletin 920).
- U.S. Training and Doctrine Command. (2024, July 24). *U.S. Army Training and Doctrine Command large language model and other generative artificial intelligence usage guidance*.
- Zhou, L., Schellaert, W., Martínez-Plumed, F., Moros-Daval, Y., Ferri, C., & Hernández-Orallo, J. (2024). Larger and more instructable language models become less reliable. *Nature*, 634, 61–68. <https://doi.org/10.1038/s41586-024-07930-y>
- Ziegler, B. (2024, May 12). The best ways to ask ChatGPT questions. *Wall Street Journal*. <https://www.wsj.com/tech/ai/chat-gpt-tips-ai-responses-d48a8f6d>.