

# Artificial Intelligence-Enabled Cyber Education: An Approach to Accelerated Education Development

By Capt. Zachary Szewczyk

I spent almost two years after I left the Cyber Protection Brigade working on training. Not traditional military training like ranges, land navigation, and vehicle maintenance, though, often to my bosses' dismay in the fledgling 3<sup>rd</sup> Multi-Domain Task Force, but rather cyber training. I wanted to teach my cyber personnel not how to handle a rifle, but rather how to handle big data; not how to read a map, but how to develop a network collection plan; not how to service a vehicle, but rather how to deploy, operate, and maintain a Security Information and Event Management system. The Army has no shortage of M4 experts, yet a worrying shortage of competent network analysts; a plethora of land navigators, yet a troubling dearth of data scientists. Yet little research has tried to answer the question, "How do we build a competent cyber workforce?" We see the consequences of this shortcoming in the news today with frequent discussions of the national cybersecurity skills gap, a problem that affects the military just as much—if not more—than the private sector. Other than vague recommendations to "start with the fundamentals", though, or "buy these seven certifications", little actionable guidance for addressing that gap exists. The fundamentals are certainly important, but what does an aspiring analyst need to learn after they understand networking? Certifications seem to answer that question—just take Network Analyst 2 after Network Analyst 1—but just punt it to someone else—and who is to say they had the right answer, or even a good one?

## Analysis of a Defensive Cyber Analyst Education Program

Little research has tried to answer the question, "How do we build a competent cyber workforce?" With few useful leads, I began to research expertise more generally. What is expertise, and how may it be defined? How can a training program facilitate the development of expertise, particularly quickly and at scale? What are the nuances of expertise in the cyber domain?

What started as a few hours of research gradually stretched into days, weeks, and then months. Thousands of pages of reading eventually led to the conclusion that rather than task mastery—the goal of training according to the U.S. Army's Field Manual 7-0: *Training* (2016)—the goal of cyber-specific training ought to be the attainment of expert-level proficiency in domain relevant areas. This is, interestingly, an important distinction that Lt. Gen. John Cushman made back in the 1970s when he advocated for *education* over training, and one with which the first commander of Training and Doctrine Command, Gen. William DePuy, strongly disagreed. (Burke) Task mastery suits static domains with well-defined tasks that are performed under a specific range of conditions and according to fixed standards—but as Cushman correctly predicted about the changing nature of warfare fifty years ago, those strictures have faded such that none of those qualifiers apply to the cyber domain today. The amorphous nature of the cyber domain demands that those operating within it cultivate both routine and adaptive expertise, the abilities to complete well-defined tasks and to solve complex problems in unfamiliar circumstances, respectively. All cyber *education*, then, should seek to develop experts—a specific term for individuals who possess both routine and adaptive expertise and are therefore capable of reliably superior performance in domain-relevant areas as a result. While no single block of instruction will ever accomplish this, all cyber education must share this common goal to make its eventual achievement a reality.

## Design of a Defensive Cyber Analyst Education Program

Drawing on operational experience and revised based on extensive research into expertise, I created a defensive cyber analyst education curriculum. This curriculum specifically focuses on developing defensive cyber analysts—a mix of host analysts who specialize in uncovering evidence of malicious activities that occur on endpoints such as user workstations and servers, and network analysts who specialize in uncovering evidence of malicious activity based on

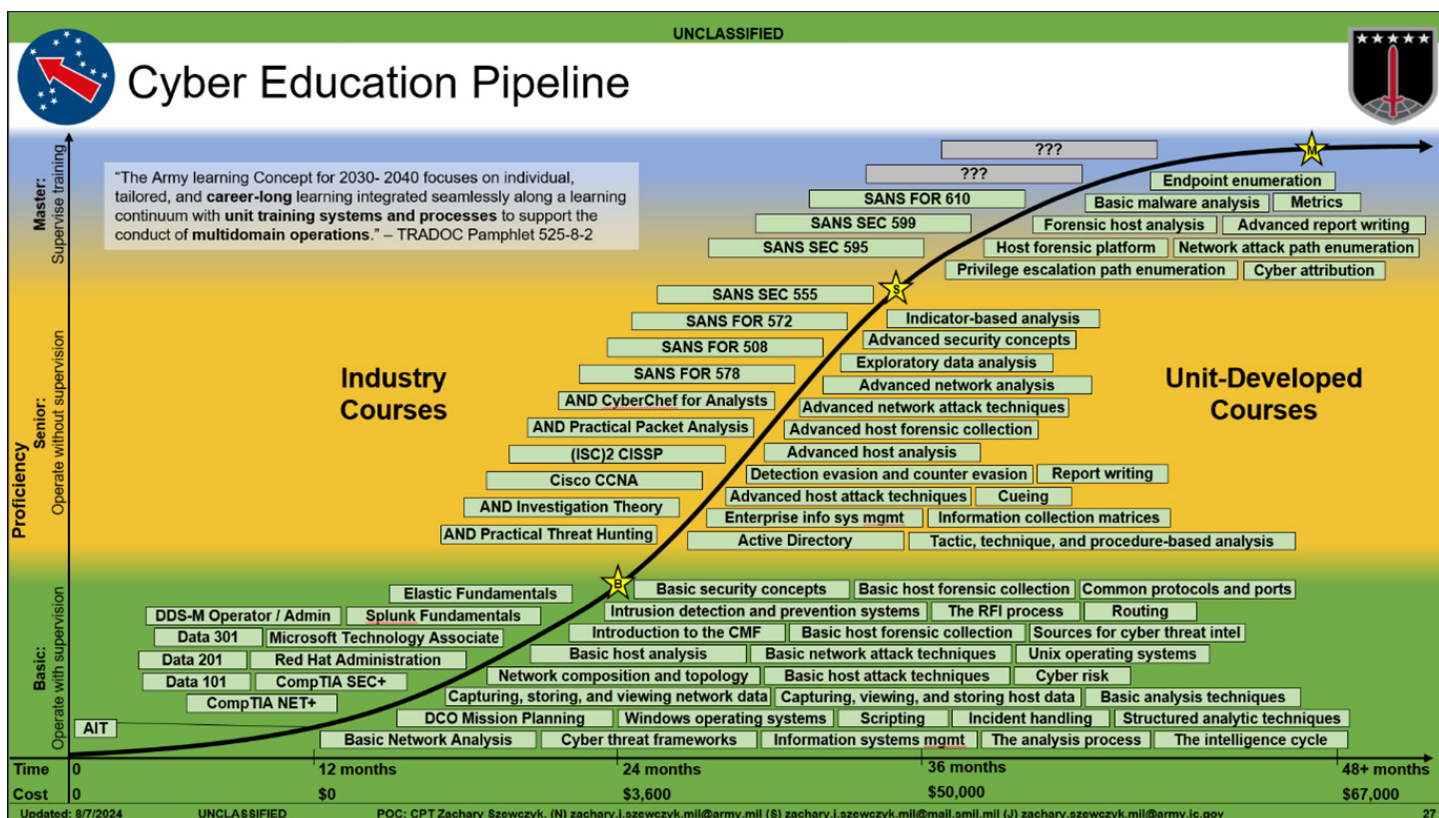


Figure 1: Defensive Cyber Analyst Education Pipeline

communications between those systems over computer networks.

Unit-developed courses, on the right side of the graph, depicts the individual lessons necessary to provide foundational knowledge and skills for defensive cyber analysts to do their jobs. At the basic level of proficiency, in the green band, these focus on developing the analysts' ability to operate under direct supervision. The corresponding industry courses, on the left side of the graph, would support that with foundational cybersecurity knowledge gained through well-known courses and certifications like CompTIA's Security+. While some have, unfortunately, begun to deride introductory-level certifications like Security+ as not worth anyone's time, I still consider these courses and their accompanying certification exams fantastic ways to establish a baseline level of knowledge and prepare individuals for higher level certifications later in their careers.

Senior-level unit-developed courses, in the yellow band, would then develop the analysts' ability to operate unsupervised and provide supervision to other, more junior analysts. The industry courses at this level would focus on work role-specific knowledge and skills through more targeted

courses like Applied Network Defense's Investigation Theory and SANS 578: *Cyber Threat Intelligence*. An emphasis on SANS's exquisite offerings will surprise no one in the cybersecurity field, but I also made it a point to consider other, less well-known but similarly high-quality courses from organizations like Applied Network Defense.

Finally, master-level unit-developed courses, in the blue band, would focus on developing the force, while the corresponding industry courses would give the analysts the deep technical knowledge to do so effectively. Many of these courses will come from SANS, at least initially, because it is rightfully considered the gold standard in cybersecurity education for a good reason. Future versions of this pipeline may feature other organizations' courses as well, such as the Naval Postgraduate School's *Data Science Certificate*.

While many other approaches to cyber education exist, mine acknowledges the critical role of internally developed courses when building a competent cyber workforce. Externally developed and hosted courses can be used to complement my curriculum, but they cannot replace it. This approach capitalizes on the Army's long-standing tradition of Soldiers training

Soldiers and avoids the pitfalls of entirely civilian-led education. While a heavy reliance on the private sector does have its merits, it is the wrong decision for the long-term health of the military's cyber forces. Operational insights are almost never available to the public, for one, and the ways, means, and ends of cybersecurity in the military—although similar to the private sector—are not the same.

### Manual Development of a Defensive Cyber Analyst Education Program

Unfortunately, a curriculum alone does not make an education program. With a plan in place, though, my small team began developing this material manually. Figure 2, below, depicts the 5-step manual instruction material development model, a product of my own design. Unlike the Army's 8-step training model, which focuses on the execution of training, my model provides guidance for creating the actual instruction material. It starts with conceptualization, then outlining, followed by shell creation, the delivery of an 80% solution, and finally the finished product at step five.

For each of the fifty-four modules on the right side of the defensive cyber analyst education pipeline, figure 1, we wrote a brief module description that consisted of a one-sentence title and a short paragraph describing the module's purpose, key topics, and a desired end state in step one. In step two, we created outlines that logically sequenced each module's topics and included a list of key points within each section. In addition to organizing the module, these outlines would also help instructors stay on track and ensure they covered key points as they taught each block of instruction. From there, we would turn that outline into actual instruction material—often a series of slides interspersed with practical exercises—that culminated in some sort of “check on learning,” such as a quiz, in steps three and four. Each module would also feature a handout with leading questions designed to enhance student engagement and facilitate guided note taking.

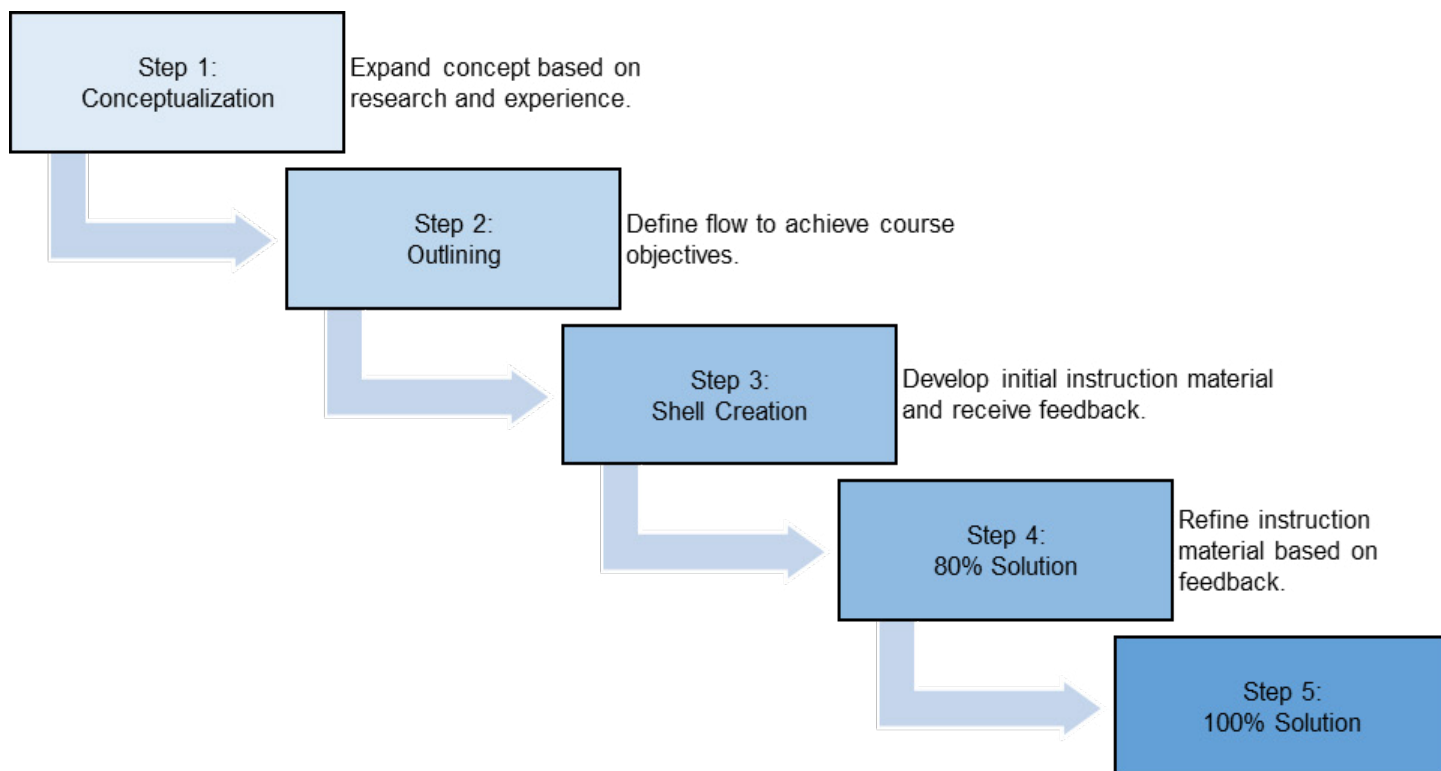


Figure 2: 5-Step Instruction Material Development Model

After an initial attempt to build all the instruction material for the entire defensive cyber analyst curriculum manually, my team estimated almost a year's worth of work to finish creating all fifty-four modules' worth of material. OpenAI's ChatGPT had grown astonishingly capable by this point, and we began exploring ways to accelerate that time-consuming development process with artificial intelligence.

## Artificial Intelligence-Enabled Development of a Defensive Cyber Analyst Education Program

Our first foray into integrating artificial intelligence into the instruction material development process had it absorb steps one, two, and three of my 5-step process: conceptualization, outlining, and shell creation. Given a module description, a large language model like OpenAI's Generative Pre-trained Transformers (GPT), via the ChatGPT web interface, would expand that description into an outline. This allowed an artificial intelligence agent to complete the initial, cursory research that fed into outlines, and attempt to logically sequence them in a coherent manner. While this seldom produced a perfect outline, it often resulted in a partial solution that one of my "course designers" could finish in short order.

Initially, this approach seemed extremely promising: in just two weeks, we used a mix of ChatGPT and Bard, a competitor to OpenAI's GPT models from Google, to create outlines for all fifty-four courses. While at first this approach seemed promising, it did not address the true limiting constraint of this process.

$$P = \frac{0.15 \times \text{Outlines} + 0.75 \times \text{Slides} + 0.10 \times \text{Handouts}}{\text{Number of Soldiers} \times \text{Number of Hours}}$$

Figure 3: Productivity Equation

Figure 3, depicts an equation I developed to measure productivity. It weighs products by the approximate amount of effort required to produce them and then calculates a rough measure of productivity as a function of products generated divided by person-hours invested to create them. When I plugged in the numbers from our first and second iterations of instruction material development, the results confirmed my suspicions: limited artificial intelligence integration had improved

our productivity over the strictly manual process, but not enough to make a significant difference.

In his 1984 book, *The Goal*, Eliyahu Goldratt introduced the theory of constraints. This theory holds that a small number of constraints—or "bottlenecks"—will limit the overall productivity of a system. In one of my favorite books, *The Phoenix Project*, Gene Kim explained this theory's applicability to business processes: "Any improvements made anywhere besides the bottleneck are an illusion. Any improvement made after the bottleneck is useless, because it will always remain starved, waiting for work from the bottleneck. And any improvements made before the bottleneck merely results in more inventory piling up at the bottleneck." In our first foray into artificial intelligence-enabled instruction material development, we had optimized for the wrong constraint!

## Artificial Intelligence-Driven Development of a Defensive Cyber Analyst Education Program

Fortunately, by then I had begun a skunkworks project to tackle the true limiting constraint: the slides. Donald Knuth's TeX typesetting language, which I used to write a guide for cyber operations called *The Handbook for Defensive Cyberspace Operations*, could also generate slides thanks to the immensely powerful Beamer package. After some tepid experimentation, I decided to dive in.

Over the course of a few hours, I developed a professional slide template in LaTeX, an extension of Knuth's typesetting language. Based on that template, a few lines of text such as those in figure 4, would now generate a PowerPoint style slide with a header, footer, unit logos, classification banner, classification markings, a title, and a bulleted list in the body.

```
\section{Title \& Content TOC Title}
\begin{frame}{Title \& Content Slide Title}
  Basic title and large, main body content
  slide.
  \begin{enumerate}
    \item Item number one.
    \item Item number two.
  \end{enumerate}
  And so on.
\end{frame}
```

Figure 4: Example LaTeX Slide Source



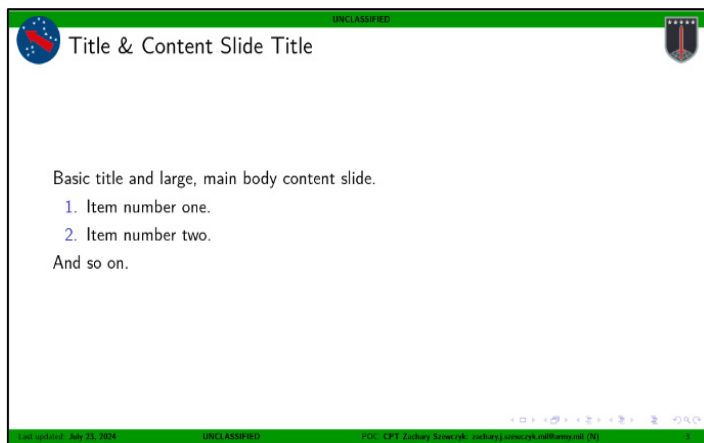


Figure 5: Example LaTeX Slide Output

Figure 5, depicts the output of figure 4's source code. By replacing everything between “`\begin{frame}`” and “`\end{frame}`” I could instead feature pictures, diagrams, flowcharts, tables—anything PowerPoint could do, I could now do with a bit of text. To call this a watershed moment in this project's development would be an understatement. Where we had once painstakingly created diagrams and tables by hand, we could now take advantage of scripting and, critically, large lan-

guage models like OpenAI's GPT4 to tackle the last true bottleneck constraining this initiative.

After a few weeks of learning to interact with OpenAI's application programming interface, or API, and developing the Python glue that would bind the entire project together, I had a working product. A series of Python scripts could now parse *The Field Guide to Defensive Cyber Analyst Education*, a short manual I wrote that explains the defensive cyber analyst education program I developed in detail, to identify all fifty-four unit-developed courses and their descriptions. The script would then feed those descriptions to OpenAI's GPT 3.5 model to generate an outline. With an outline and a series of related course objectives, the more capable GPT4 model would revise the outline into a more detailed, finished product. GPT4 would also create the handout to accompany the course material. These steps alone underscore the immense power of *generative* pre-trained models, which accepted just under 5,000 words as input and output over 60,000 words in outlines and handouts.

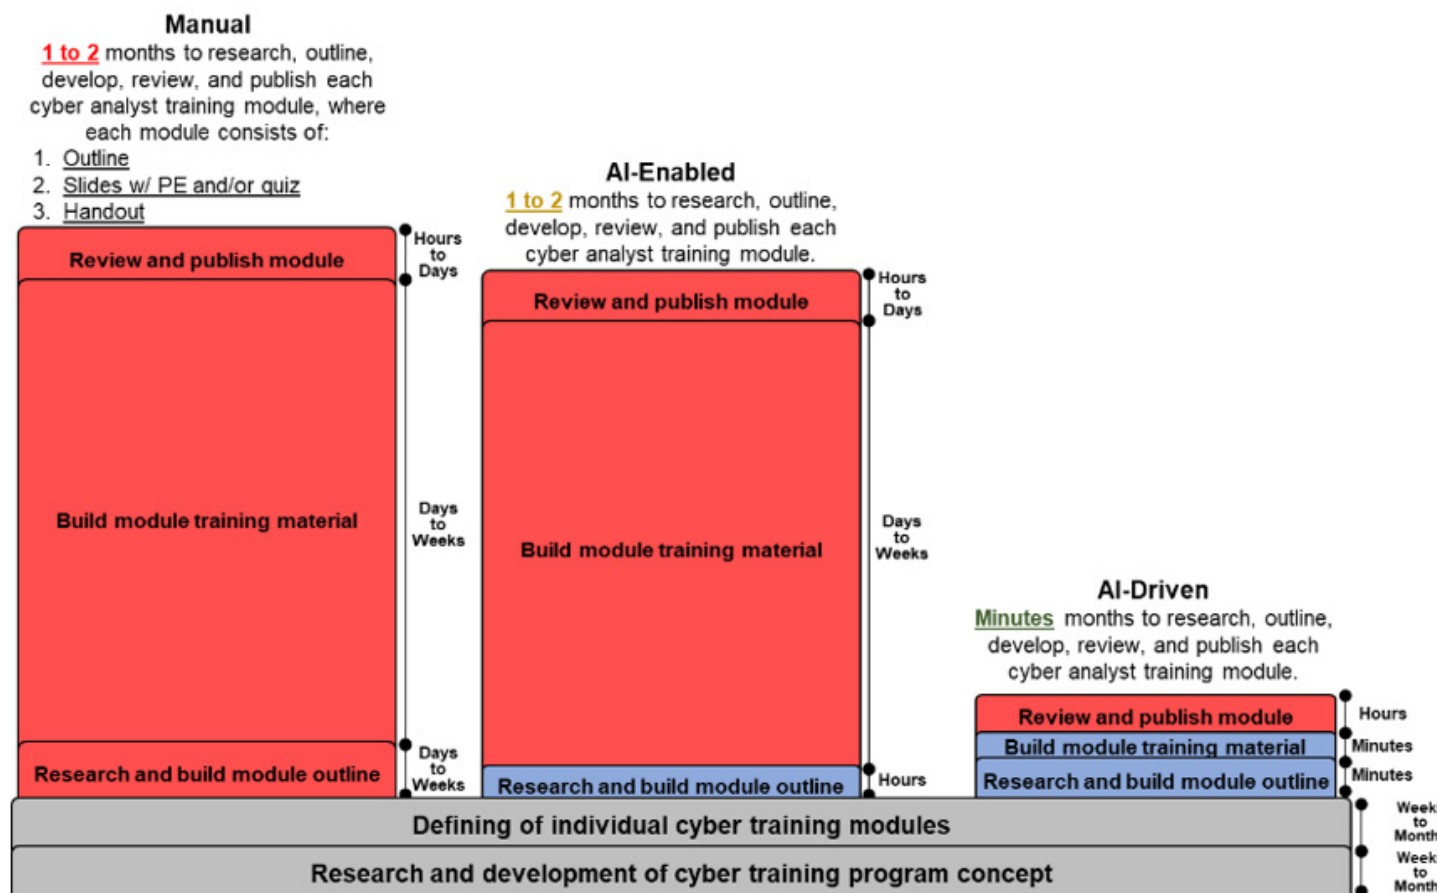


Figure 6: Manual vs. AI-Enabled vs. AI-Driven Instruction Material Development Process

Finally, the script would read these outlines and iteratively prompt GPT3.5 and GPT4 to generate individual slides. These slides would then be stitched together into a complete presentation using another extension of Knuth's TeX, called XeLaTeX, via the XeTeX engine. Here, the original 5,000 words of module descriptions became 60,000 words in outlines and handouts, which expanded into a staggering 284,000 words on 1,600 slides across 54 presentations in class material. Through scripting and the help of artificial intelligence, we had successfully automated the entire 5-step instruction material development model. Figure 7, compares the three incarnations of the instruction material development process by the approximate amount of time necessary to complete each step: manual, AI-enabled, and AI-driven. What would have taken months under the best of circumstances if done the old, manual way took mere seconds and cost me just \$34.68.

Aside from speed, this programmatic, AI-driven approach to content generation also had another benefit: machine-readable data structures and interfaces made transforming content a few minutes' work with a Python script. In addition to generating 54 individual slide decks, this pipeline also generated an accompanying book for each module. Each book contained the same material as the original course content, for those more inclined to learn through reading than by listening to a lecture.


This approach also had other benefits from an administrative perspective, too. For example, compiling all the slides and books into a single document for review by a foreign disclosure officer took a few seconds rather than hours of copying-and-pasting hundreds of slides into "Master PowerPoint v7.ppt". Condensing the outlines into a nice catalog for dissemination to other organizations required a few lines of Python, not hours wrestling with Microsoft Word.

By focusing on and optimizing the correct constraint, I created a process that took months of work and reduced it to just a matter of hours. Figure 8, compares the productivity measures for the three approaches.

$$P_{\text{manual}} = \frac{0.15(15) + 0.75(15) + 0.10(2)}{10 \times 480}$$

$$P_{\text{manual}} = \frac{2.25 + 11.25 + 0.20}{4800}$$


$$P_{\text{manual}} = \frac{13.7}{4800}$$

$$P_{\text{manual}} = 0.002854$$


$$P_{\text{AI-enabled}} = \frac{0.15(54) + 0.75(1) + 0.10(54)}{2 \times 120}$$

$$P_{\text{AI-enabled}} = \frac{8.1 + 0.75 + 5.4}{240}$$

$$P_{\text{AI-enabled}} = \frac{14.25}{240}$$

$$P_{\text{AI-enabled}} = 0.059375$$


$$P_{\text{AI-driven}} = \frac{0.15(54) + 0.75(54) + 0.10(54)}{1 \times 24}$$

$$P_{\text{AI-driven}} = \frac{8.1 + 40.5 + 5.4}{24}$$

$$P_{\text{AI-driven}} = \frac{54}{24}$$

$$P_{\text{AI-driven}} = 2.25$$

*Figure 7: Manual vs AI-Enabled vs AI-Driven Instruction Material Development Process Productivity*

Artificial intelligence tools like OpenAI's ChatGPT have taken the world by storm. Their sudden popularity, and the accompanying "AI-ification of everything", makes it easy to forget that this technology is still in its infancy. Many organizations, including the Department of Defense, are still exploring appropriate roles for it, and trying to understand its impact. As I look back on the first phase of this project, I have answers to both of those questions, and the results to back them up. Instruction material generation is a fantastic role for AI, particularly when paired with domain experts and used in an iterative manner. I know, because it ultimately led to an 788x increase in our productivity.

## Way Forward

As I look back on this project, and the months of research that enabled that execution to succeed, I am immensely proud of how far this initiative has come. I am also excited for the future as I consider all the opportunities to improve and expand this cyber education program.

The current incarnation of this program focuses on U.S. Cyber Command's Host Analyst and Network Analyst work roles. Given the continued difficulty of effective intelligence support to cyber operations, I look forward to expanding its scope to include a cyber threat intelligence analyst capacity as a small step toward remediating that. In a similar vein, I also look forward to exploring what it means to train officers and NCOs in the now-defunct Cyber Network Defense (CND) Manager work role, which the Army unfortunately nixed several years ago. Planning, overseeing, and executing defensive cyber operations has become a responsibility shared by the Cyber Planner and Analytic Support Officer work roles, but I have and will continue to advocate for an important third leg to this stool, the CND Manager, who handles the day-to-day execution of cyber operations, leads analysis, and coordinates incident responses. Fortunately, integrating courses to build cyber threat intelligence analyst and cyber network defense manager capacities will result in a logarithmic increase, not a linear one, thanks to the integrated nature of this program. By designing this program around knowledge domains rather than work roles, adding sufficient materials will require minor course adjustments instead of drastic changes in direction.

I believe this approach has the potential to apply elsewhere as well. Applying a similar artificial intelligence pipeline to areas sorely in need of formal curriculum, such as the electronic warfare specialty, could help grow this nascent field.

Unfortunately, generalizing this pipeline to other work roles—and even other fields—is not without risk. Accelerating the instruction material development process risks flooding the space with low-quality products. Appropriate direction, important now to economize resources, will become critical in a future free of such constraints. Outcome-based learning is the right approach,

particularly for cyber where Soldiers must be educated not trained, but the outcomes achieved must become job qualification. Knowledge for knowledge's sake is the purview of academia, not the military.

General-purpose models like GPTs 3.5 and 4, although effective for developing defensive cyber analyst training given the field's significant overlap with cybersecurity in the private sector, are also unlikely to perform well in narrow specialties throughout the military. Fortunately, phenomenal initiatives like CamoGPT will soon provide Soldiers with access to large language models trained on domain-specific information and backed by military doctrine. CamoGPT must, however, be appropriately resourced to support state of the art, frontier models. Many "large" language models, with just a few billion parameters, hardly deserve the name compared to those with trillions of parameters available today. Emergent properties, especially important in ill-structured tasks like training development, do not appear in small models, and only begin to appear in some of the largest models available today. CamoGPT must have the resources to handle these gargantuan models lest it become little more than a toy.

## Conclusion

This article's approach represents one of the few attempts to codify and disseminate a formal approach to cyber analyst education, particularly one that views internally developed courses as central to its execution rather than an afterthought. I hope to see other units in the cyber mission force seize this opportunity to collaborate and build upon this program. The Army has no shortage of M4 experts, yet a worrying shortage of competent analysts, and while this program may not be *the* answer, it is certainly a great start.

### **About the Author**

Captain Zachary Szewczyk commissioned into the Cyber Corps in 2018 after graduating from Youngstown State University with an undergraduate degree in computer science and information systems. He has supported or led defensive cyberspace operations from the tactical to the strategic level, including several high-level incident responses. He has served in the Cyber Protection Brigade and the 3rd Multi-Domain Task Force.

### **References**

Burke, E. M. (n.d.). Ignoring Failure: General DePuy and the Dangers of Interwar Escapism. *Military Review*, 42–58.