Teaching Generative AI for Cybersecurity: A Project-Based Learning Approach

Nate Mathews
Department of Cybersecurity
Rochester Institute of Technology
Rochester, NY, United States
nate.mathews@mail.rit.edu
0000-0001-6186-7001

Christopher Schwartz
Department of Cybersecurity
Rochester Institute of Technology
Rochester, NY, United States
ccsics@rit.edu
0000-0001-6867-4202

Matthew Wright
Department of Cybersecurity
Rochester Institute of Technology
Rochester, NY, United States
Matthew.Wright@rit.edu
0000-0002-8489-6347

Abstract—In the Spring 2024 semester, we introduced an elective course titled "Generative AI and Cybersecurity" for MS and upper-division BS students specializing in cybersecurity at our university. The course was designed to equip students with a foundational understanding of Generative AI, particularly large language models (LLMs) like GPT-4, and explore their applications within the field of cybersecurity. Through a combination of classroom instruction, hands-on projects, and industry guest lectures, students engaged with the technical, ethical, and legal dimensions of AI in cybersecurity. The course emphasized practical learning, with students gaining experience in AI tools such as ChatGPT, as well as developing skills in prompt engineering and API usage. While some students were eager for even more technical AI content, they appreciated the hands-on learning, insights from industry guest speakers, and the chance to see how the more powerful models like GPT-4 could be usefully applied to cybersecurity problems.

Keywords—*Generative AI, Cybersecurity, Large language models, AI in education*

I. INTRODUCTION

The rapid evolution of artificial intelligence (AI), particularly in the realm of large language models (LLMs) such as GPT-4, has begun to transform various industries. Cybersecurity is confronted with both immense opportunities and significant challenges from AI, from automating threat detection to raising new ethical and legal concerns. Recognizing this paradigm shift, we introduced an elective course for MS and upper-division BS students entitled "Generative AI and Cybersecurity" during the Spring 2024 semester.

We had two main motivations: to equip students with a deep, practical, and ethical understanding of generative AI technologies, and to explore how these tools can be applied effectively in cybersecurity contexts. We designed the course to take a hands-on, project-based approach, thereby enabling students to interact directly with cutting-edge generative AI tools. By integrating prompt engineering, *retrieval-augmented generation* (RAG), and AI-assisted code development into the coursework, we gave students opportunities to apply AI in practical cybersecurity tasks. Ethical and legal discussions around AI misuse, bias, and privacy were also central to the course, encouraging students to reflect on the broader societal impacts of these technologies.

Terminology Note. AI has been used in cybersecurity for some time, but the rise of *generative* AI introduces new challenges. Generative AI leverages neural networks, such as large language models (LLMs) and diffusion models, to create high-quality textual, audio-visual, and even code-based content. In addition to content generation, these neural networks excel in pattern recognition, making generative AI applicable beyond content creation. In this paper, unless otherwise specified, the term "AI" will refer specifically to generative AI.

Key Findings. This first iteration of the course yielded several key findings:

- Students gained hands-on experience with AI tools, increasing their confidence in applying generative AI to cybersecurity tasks.
- Ethical discussions broadened students' perspectives on AI's societal implications in cybersecurity, but some students felt that more should have been done.
- Feedback indicated a need for more in-depth technical explanations and clearer assignment structures.
- Students requested stronger connections between AI concepts and real-world cybersecurity applications like penetration testing and malware analysis.

These insights will guide future iterations of the course. Of particular concern is refining content and strengthening the connection between AI tools and practical cybersecurity tasks.

The course syllabus, lecture slides, and assignment materials are available online via the CLARK cybersecurity library.

AI Disclosure. In the spirit of this project, we used Generative AI heavily in the development and execution of all stages of this work, including in the writing of this article.

II. LITERATURE REVIEW

As noted at the outset, generative AI has recently emerged as a powerful tool across many fields, including cybersecurity. These models can assist with tasks like threat detection, vulnerability assessment, and security policy automation by analyzing vast amounts of data and generating actionable insights [1]–[3]. However, alongside these advancements, generative AI introduces challenges, such as the potential for biased outputs, a lack of transparency in decision-making, and vulnerabilities to misuse by malicious actors [4]–[7].

A. Generative AI in Education

The exploration of AI's role in education has become popular recently, revealing both significant opportunities and challenges across various disciplines [8]–[12]. AI-driven tools have the potential to enhance personalized learning experiences and provide real-time feedback to students [13], [14]. For instance, AI can approximate a one-to-one teacher-to-student ratio by offering individualized feedback on coding practices [13], [14]. Researchers have also emphasized the importance of balancing the benefits of AI with risks such as over-reliance and academic misconduct [11], [12], [15]. Concerns about ethical considerations, biases, and data privacy are also prevalent [8], [16], [17], highlighting educators' need to incorporate critical thinking and ethical reasoning into the curriculum. Addressing ethical challenges such as misinformation and over-reliance is crucial to avoid hindering essential cognitive processes in education [9], [12].

In cybersecurity education, integrating AI tools presents unique opportunities and challenges. Balancing AI/ML and cybersecurity knowledge is essential to enhance learners' abilities in real-world applications like malware detection [18]. Using large language models can streamline curriculum development, aligning educational content with industry standards such as OWASP [19]. Frameworks can guide ethical considerations and help tailor AI integration to support personalized learning without compromising critical thinking and ethics [20], [21]. Educators must foster critical thinking and self-regulated learning strategies [22] and prepare students for effective use of AI tools through skills like prompt engineering [17].

III. COURSE DESIGN

The design of the "Generative AI in Cybersecurity" course moved away from traditional lectures, and instead focused on experiential learning, collaboration, and critical reflection.

Structure. The course was structured as a 15-week elective, with approximately 25 hours of classroom instruction augmented by student presentations of their three projects and several guest lectures from speakers with real-world experience in industry. Table I presents a general overview of the time commitment for each course component. The main topics covered by the course were the following:

- An introduction to Generative AI models and their functionalities.
- Technical aspects and practical usage of AI, including API calls, locally-hosted models, prompt engineering, RAG, and fine-tuning.
- Applications of Generative AI in cybersecurity.
- Ethical implications of Generative AI in terms of privacy, bias, and accountability.
- Regulatory compliance and the intersection of AI with data protection laws.
- Case studies driven by student interests, covering areas like incident response, cyber threat intelligence, and red teaming operations.

The course was organized so that students started by completing simpler text-generation tasks before advancing to more sophisticated challenges such as AI-assisted code development and RAG applications. Sessions on AI ethics and legal compliance were interspersed throughout the course, and time slots were made to accommodate invited speakers.

TABLE I. Course Timeline by In-Class Days

Component	Days Spent
Introduction to Generative AI: Writing, Coding, and Tools	6
Ethics and Legal Implications of Generative AI	4
Advanced Techniques: Fine-tuning, Retrieval-Augmented Generation, and Specialized Tools	5
In-class Student Project Development and Presentations	8
Invited Guest Lectures from Industry and Government Experts	5

Projects. Central to the course were three major team-based projects, which serve as the primary mechanism for student learning. Each project introduced new generative AI capabilities within a cybersecurity context:

1) Project 1: Generating Cybersecurity Documentation

Students used prompt engineering techniques to generate cybersecurity documents, such as threat assessments, policy guidelines, or incident response protocols. The goal was to explore how AI can be employed in document generation while recognizing its limitations. Particular attention was paid to the accuracy and relevance of the generated content. To achieve this, students were tasked with defining clear objectives and identifying target audiences for their documents. They utilized generative AI to craft initial drafts, which they refined through iterative evaluation and integration of human expertise. This process emphasized understanding key cybersecurity concepts and tailoring content to specific use cases, such as simulating scenarios for IT administrators or business executives. The final deliverables included a polished

document and a presentation outlining the AI's role and the lessons learned throughout the project.

2) Project 2: AI-Assisted Code Development

Students applied generative AI tools to develop non-trivial but basic software solutions relevant to well-established cybersecurity tasks. Successful proofs-of-concept were not a priority; rather, the goal was for students to explore the *potential* of AI for code generation and the iterative nature of refining AI-generated outputs. Teams selected application areas such as network analysis or authentication, and leveraged AI to brainstorm ideas, generate code modules, and design testing frameworks. Students crafting effective prompts generated initial code components, which were iteratively refined and tested. The project also directed the creation of a functional demo to showcase key features. This experience highlighted the benefits and challenges of integrating AI into software development, including debugging AI-generated code and balancing human input with automated suggestions.

3) Project 3: Advanced Generative AI in Cybersecurity.

Students employed RAG techniques or other advanced generative AI capabilities to develop innovative cybersecurity applications. The project began with refining ideas based on feedback from earlier proposals and selecting appropriate tools and data sources. Teams implemented core components of their projects, such as using AI to simulate integration scenarios or generate test cases. This stage emphasized the practical application of generative AI to address complex cybersecurity problems, such as automated threat detection or data enrichment for analysis. We allowed teams to develop new ideas or build upon their application from Project 2. Students documented their processes and outcomes in a comprehensive report and reflected on team dynamics and the limitations of AI in their solutions. The project culminated in a presentation in the last week of class, allowing teams to share their implementations and insights with peers.

In the Appendix, we present samples of the student-produced projects for these projects in Figures 1, 2, and 3. Additional materials from these project samples are available alongside our course material on CLARK.

Tools. Throughout the projects, students used a variety of AI tools. For the course duration, the students were given paid access to OpenAI's ChatGPT and the OpenAI API [23]. The students were also encouraged to explore alternative tools, such as the free versions of Microsoft CoPilot [24] and Anthropic's Claude [25]. Most students used the AI-powered Gamma to build their presentation slides [26], while the instructors also used Tome [27]. We also encouraged and saw wide use of the cloud-based Google CoLab [28] to enable collaborative code development and access to GPU resources.

Addressing Privacy and Ethical Concerns. At the beginning of the course, we held an open discussion to ensure that students were aware of the privacy implications of uploading documents, code, or other resources to online LLM applications. This discussion emphasized responsible and ethical usage of AI tools, particularly regarding sensitive or proprietary information. For all assigned projects, students were expected to develop their work from scratch rather than uploading or modifying existing code bases. However, we lack enforcement mechanisms to ensure that students are adhering to ethical practices when handling data and code resources. Furthermore, we used the paid version of ChatGPT, which provides an option to disable OpenAI's use of user data for training purposes. This safeguard helped mitigate concerns about inadvertent data sharing and served as a model for the responsible integration of AI into the learning process.

Trending Topics. A useful component of the course was a regular news presentation assignment, designed to keep students informed about current developments in generative AI and cybersecurity. Each week, students were tasked with selecting recent news articles or research papers on the intersection of AI and cybersecurity. They were required to present a five-minute presentation of the article and discuss aspects of the content that they found interesting.

Assessment of Learning Outcomes. A unique feature of this course was its "ungrading" system, which prioritized students' self-assessments over instructor evaluations [29]. After each project, students submitted self-reflection essays evaluating their performance, contributions, and learning outcomes. Prompts embedded in the assignments guided students to reflect on aspects such as the impact of generative AI on team dynamics and project planning, their personal contributions, and lessons learned. On the basis of these reflections, instructors provided extensive feedback in order to guide the students toward balanced and honest assessments.

A key feature of ungrading is student self-assigned grading: students submitted a final reflective essay of their experience, including an overall grade for themselves at the end of the semester. The instructors retained the right to adjust these grades if they fell far outside observed performance and peer reports about their work. While we anticipated the need for one-on-one meetings with students whose self-assigned grades deviated from our evaluations, this was ultimately unnecessary. Nearly all of the self-assigned grades were found to be reasonable and consistent with our observations, highlighting the effectiveness of the structured guidance provided to students during the self-assessment process. We would, however, have widened the grading scale overall, lowering roughly half of the A's to A's, for example.

Prerequisites and Student Context. Enrollment in the course required students to have completed an intermediate-level cybersecurity course and to have participated in at least one cooperative education (co-op) experience – typically a paid internship in cybersecurity. These prerequisites ensured that students possessed foundational cybersecurity knowledge as

well as some exposure to professional environments, enabling them to critically evaluate the utility of generative AI tools in real-world scenarios. Notably, no prior experience with artificial intelligence or machine learning was required, allowing students to engage with these concepts from the ground up. The course was attended by 28 students, comprising advanced undergraduate and graduate students, with diverse technical backgrounds within cybersecurity.

IV. GOALS AND INTENDED OUTCOMES

Our primary goal as instructors was to equip students with both the technical proficiency and critical thinking skills necessary to navigate the complex intersection of generative AI and cybersecurity. We aimed to enable them to apply AI practically in real-world teamwork contexts while considering the ethical and legal implications of these applications.

Course Goals. The following overarching goals guided the design and implementation of the course:

1) *Basics of Generative AI*. First, we aimed to provide students with a foundational understanding of how generative AI functions. This included gaining proficiency in using generative AI tools such as ChatGPT and locally-hosted models, understanding their underlying architecture, and mastering techniques like prompt engineering, retrieval-augmented generation (RAG), and model fine-tuning.
2) *Practical Application*. Second, we sought to explore the practical applications, benefits, and limitations of integrating generative AI tools within various cybersecurity tasks. Through hands-on projects that simulated real-world scenarios—such as generating threat assessments, AI-assisted code development, and incident response planning—students applied AI tools effectively and critically evaluated these techniques in real-world cybersecurity applications.
3) *Teamwork*. We aimed to foster collaborative skills and teamwork. Students engaged in group projects that emphasized collaborative problem-solving, enhancing their communication, leadership, and ability to work effectively within a team—key skills in both cybersecurity and AI-driven fields.
4) *Ethics and Law*. We encouraged students to engage in thoughtful ethical and legal discussions pertinent to the use of AI in cybersecurity. They reflected on how AI could impact decision-making in security contexts, identified ethical issues and risks such as data privacy, bias in AI algorithms, potential AI misuse, and the accountability of AI-driven systems, and considered how AI could adhere to or complicate regulatory compliance.

Linking Goals to Evaluation. The course's goals and outcomes were not only integrated into the instructional content but also embedded within the evaluation framework. Students were assessed based on their ability to demonstrate technical proficiency with AI tools through projects that applied generative AI in solving cybersecurity problems; collaborate effectively within teams, showcasing strong project management and communication skills; critically engage with the ethical and societal issues surrounding AI in cybersecurity, both in their reflective reports and during class discussions; and provide thoughtful analysis of the limitations and potential misuse of AI in cybersecurity, reflecting an awareness of broader regulatory and legal contexts.

Pedagogical Approaches. Our pedagogical choices for this course emphasized *active learning and practical engagement*, focusing on experiential education designed to bridge theoretical concepts and real-world applications. Although we regularly deployed traditional lectures, given by both ourselves and guest speakers, project-based learning (PBL) was the heart of the course.

Each project aligned with the typical PBL cycle: framing real-world scenarios, fostering active exploration, and culminating in reflective practice. For instance, students tackled practical challenges such as generating threat assessment documents, building AI-assisted tools, and integrating generative AI into incident response workflows. These projects began with analyzing the problem context and defining objectives, guided by instructors to ensure relevance to real-world cybersecurity applications. While the short timeline for each project (2–3 weeks) limited formal reflection during the process, students were encouraged to engage in frequent team discussions to address challenges and refine their approaches as the project progressed.

We chose PBL as the primary instructional method, aligning with learning theories that emphasize the importance of students actively constructing their own knowledge [30]. In cybersecurity, where the landscape evolves rapidly, traditional passive learning approaches may fail to engage students with the dynamic nature of real-world challenges. Each project challenged students to identify a meaningful real-world scenario, facilitating *experiential learning* and allowing students to gain hands-on experience with the tools and processes they are likely to encounter in the field. The *iterative process* embedded in the projects—in which students refined their work through feedback and reflection—mirrored industry workflows, promoting *continuous improvement* and real-world problem-solving skills.

Collaboration was also central to our pedagogical strategy. Students worked in teams on all projects, a design choice rooted in the idea that *collaborative learning* enhances individual and collective problem-solving [31]. Group work allowed students to bring different perspectives and expertise to the table. They were encouraged to engage in peer feedback and discussion throughout the course, helping develop critical *teamwork skills*, which are essential in professional cybersecurity.

Reflective learning was key to the course, with students completing *individual reflection reports* after each project to assess their learning, contributions, and the role of generative AI. This was reinforced by the ungrading approach, where

students evaluated their own effort, outcomes, and growth. Such reflective exercises are crucial for developing *metacognitive skills*, enabling students to not only understand the content but also monitor and improve their own learning strategies.

We also invited guest speakers from the cybersecurity industry who provided real-time insights into current trends, challenges, and innovations in applying generative AI to cybersecurity. By hearing from professionals actively working in the field, students could contextualize their projects within broader industry practices and emerging technologies. The inclusion of guest speakers served to bridge the gap between academia and industry, giving students a clearer sense of how the skills they were developing would translate to their future careers.

Another crucial aspect of our pedagogical approach was encouraging students to critique the AI tools they were using. Beyond learning how to use AI for cybersecurity tasks, students were prompted to analyze the limitations, biases, and ethical concerns associated with these tools. This critical engagement is vital in preparing students to navigate the complex ethical landscape they will face as professionals.

V. STUDENT OUTCOMES

To evaluate the effectiveness of the course and assess student outcomes, we used several methods. First, with approval from our university's IRB, we conducted a post-course survey to gather quantitative and qualitative feedback from 17 student respondents. Specifically, we asked them to assess the extent to which they perceived their knowledge and competence in the practical and ethical components of applying generative AI in cybersecurity had grown due to the course. They provided their assessments using both numeric scales and written answers. The results of this post-course survey are presented in Table II. Second, most students completed a course evaluation form following standard university procedures. These evaluations offered insights into how the students perceived the course's overall structure and supplied additional perspectives on its effectiveness and areas in need of improvement. Third, self-reflections written by the students throughout the course provided additional insights into their individual learning processes, challenges faced, and overall experience.

These materials were then cross-referenced with the student's outputs, specifically the students' project and news presentations, class participation, and their self-assigned grades along with the reasoning they provided for these grades.

Student Views. Overall, students expressed a high level of satisfaction with the hands-on, project-based approach and the integration of generative AI into cybersecurity tasks. They particularly valued the opportunity to engage in real-world applications of generative AI, emphasizing how the projects helped them understand the impact of AI in cybersecurity. Including guest speakers and interactive group projects was also well received; students appreciated hearing from industry professionals. One student stated, *"I found the guest lectures particularly useful, especially Jason Ross, who was able to talk about the problems he is working on with AI at his workplace."* Additionally, students appreciated the flexibility in how they approached their projects, with access to cutting-edge AI tools. One participant commented, *"A tool that I found particularly useful was having access to ChatGPT 4. GPT-4 is more advanced than the publicly available 3 and 3.5. Additionally, it provides access to image generation."* We noticed that there was a significant gap in capabilities between the paid and free generative AI services, and making the most capable agents available to students was key to reducing friction and frustration.

While the feedback was generally positive, students also identified several areas for improvement. Some expressed a need for more concise assignment instructions. We note that the assignments were generated with a lot of help from ChatGPT, which tended to provide wordy outputs if not guided to be more concise. Although guest lectures were highly appreciated, some students felt that more attention should be placed on core course content. One remarked, *"I wish that we had spent more time on RAG, and on more practical methods for commercial deployment."* Several students found certain topics, such as prompt engineering and AI model mechanics, challenging. As one reflected, *"Generative AI simply isn't ideal for certain things, like coding complex cybersecurity tools. We often overcame these issues by repeatedly iterating on code indefinitely until it functioned in a way we could be satisfied with."* They requested more individual learning opportunities and in-depth technical explanations. Some wished for more examples of how generative AI could be directly applied within cybersecurity. One participant noted, *"Doing actual pen-testing, malware analysis, network analysis with Gen AI—but under a guided hand so we understand how it works rather than individual projects."* This weakness was also evident from the quantitative scoring, where student confidence in applying generative AI "to address cybersecurity problems" ranked the lowest among the skills surveyed. Importantly, the students were roughly divided between those who felt there could have been more exploration of the ethical and legal challenges of AI in cybersecurity, and those who felt what was explored was sufficient. For example, one student felt *"we could have done an ethics assignment, maybe a poster or short paper (individually),"* while another stated, *"the course covered a wide range of topics and discussed ethics extensively."*

TABLE II. A reporting of the quantitative results of an optional post-course survey provided to the students

Students were asked to provide a rating in the range of 1 (low) and 5 (high), or a binary positive/negative response for each question. The questions are presented verbatim, and the scale instructions are removed.

1	How would you rate the extent to which the course enhanced your knowledge of the following topics concerning Generative AI?	Average Score
	• Technical aspects	3.81
	• Impacts on cybersecurity	3.81
	• Implications for ethics and law	4.00

2	How would you rate your confidence in your ability to do the following before and after this course?	Before	After
	• Use Gen AI to accelerate the speed of document creation?	2.83	4.39
	• Use Gen AI to improve the quality of document creation?	2.61	4.17
	• Use Gen AI to increase your productivity in code development?	2.83	4.17
	• Use Gen AI to extend your coding skills and capabilities?	2.89	4.11
	• Use Gen AI to address cybersecurity problems?	2.11	3.50
	• Use Gen AI to develop cybersecurity tools with new functionalities?	2.39	3.78

3	How would you have rated the likelihood you would have used Generative AI in your future cybersecurity work before you took this course, and will use it after having taken it?	Before 3.06	After 4.17

4	How would you rate the effectiveness of the hands-on projects in the following areas?	Average Score
	• Increasing your practical knowledge of Gen AI	4.17
	• Increasing your practical knowledge of Gen AI specifically for cybersecurity	4.06
	• Increasing your ethical considerations about applying Gen AI	3.50

5	Based on your experience in the projects, how would you assess your experience of Generative AI's impact on teamwork?	Positive	Negative
	• Brainstorming	94%	6%
	• General coordination	47%	53%
	• Distributing workload	47%	53%

Challenges. During the course, several technical and logistical challenges emerged. Many students found AI models daunting, particularly with topics like RAG. A lack of examples of successful applications of generative AI left some feeling directionless at times. One student explained, *"A major problem I had was collaborating on specific tools when there weren't collaborative tools already implemented for the product. And we got around this by picking out specific AI chats and defining prompting for it."* Students frequently encountered difficulties when debugging AI-generated code, especially when the generative AI produced inconsistent outputs. Another noted, *"Trying to debug code that was generated by ChatGPT using libraries that I was unfamiliar with was a challenge since I wasn't really able to get it to correct its own mistakes and I wasn't familiar enough with the libraries to do it myself."* While AI tools were integrated into project work, some students felt they needed more guidance on practical cybersecurity-specific applications. One remarked, *"I felt that I had to learn a lot about real-world deployment outside of class."*

Evaluation and Impact. Our analysis of student surveys and project outcomes revealed several key insights into the course's effectiveness. Quantitative assessments showed substantial gains in students' understanding of both technical and ethical aspects of generative AI. Students reported significant improvements in their knowledge of AI tools and

techniques, with average ratings of technical understanding improving from moderate to high across the board. One student reflected, *"I really enjoyed learning about Claude. ChatGPT gets most of the news coverage, and while impressive, I appreciated the way Claude is a bit more geared towards writing."* The hands-on projects were consistently rated as highly effective, and students felt that the experiential learning approach enhanced their practical abilities. *"Using ChatGPT to code was an interesting tool I've never used before. While it seems like a problem to rely on, it seems to at least be good for certain shortcuts,"* one student noted. The course positively influenced students' views on the use of generative AI in their future cybersecurity careers. The likelihood of using AI tools in professional settings significantly increased, as expressed by a student: *"Using RAG to optimize output was really interesting, since the barriers of entry were low enough that I could see it being practical for me to apply in my career."* Although students were divided on the amount of ethical and legal exploration in the course, they were generally satisfied with its quality. Many reported being strongly impacted by the in-class ethical discussions and reflections, gaining a deeper awareness of AI's societal implications. One shared, *"Discussing how [deepfakes and model biases] can 'harm' in a non-physical way was a very interesting discussion and put emphasis on why we need more laws and security in Gen AI."*

VI. CONCLUSION

Our university course "Generative AI and Cybersecurity" largely succeeded in achieving its primary goals, particularly in the areas of hands-on skill development, theoretical understanding, and ethical awareness. However, the students also consistently indicated that they felt there was a need for improvements in the following aspects of the course:

- More detailed and structured technical content, particularly on topics like RAG and AI model fine-tuning.
- Enhanced support resources, including tutorials and guided exercises for complex tools.
- A clearer focus on how AI tools can be applied to specific cybersecurity tasks, such as penetration testing, malware analysis, and network defense.

These improvements can certainly be enacted in future iterations of the course. We admit that, as instructors, our own knowledge of how AI is applied in cybersecurity was limited due to the fast-changing nature of this topic. We had to rely ourselves on online resources and talks from experts. So building that knowledge will be essential to offering this course in a way that maximizes student learning. Nevertheless, given our limitations as instructors, we were able to leverage PBL to get students to engage semi-independently with the tools and reach their own understanding of the possibilities of the technology to date.

In addition to technical and resource-focused adjustments, addressing student concerns regarding the clarity of assignment instructions will be an important area of improvement. Many of the assignment prompts were heavily influenced by generative AI, which, while reducing preparation time, occasionally led to verbose or ambiguous instructions that challenged student interpretation. To resolve this, future iterations will incorporate a more iterative process for testing and refining assignment instructions, potentially involving pilot reviews by peer or graduate assistants. We also aim to integrate scaffolding techniques to break assignments into smaller, clearly defined steps and provide concrete examples or templates for reference. Furthermore, incorporating periodic student feedback mechanisms during the course will help identify and address misunderstandings in real time.

More challenging is the precise role for ethical-legal exploration going forward. As noted earlier, our students were split about this component of the course. While they were all satisfied with the quality of ethical-legal exploration, they were divided about its amount, with some wanting more and some feeling that what was done was sufficient.

One strategy to address the perceived need for more would be to create a ethics and law module in future iterations of the course, in which ethical-legal exploration would be given three straight weeks of focus and a dedicated assignment. Another approach might be to add an explicit ethics and law sub-assignment to each major project, so that it is addressed regularly throughout the coursework. Exploring these options will be interesting in future work.

ACKNOWLEDGEMENT

This project was sponsored by the National Security Agency under Grant Number H98230-21-1-0166. The United States Government is authorized to reproduce and distribute reprints notwithstanding any copyright notation herein.

Disclaimer - Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Security Agency.

REFERENCES

[1] D. Saha, S. Tarek, K. Yahyaei, S. K. Saha, J. Zhou, M. Tehranipoor, and F. Farahmandi, "Llm for soc security: A paradigm shift," *IEEE Access*, pp. 1–1, 2024. [Online]. Available: https://doi.org/10.1109/ACCESS.2024.3427369

[2] T. Ali and P. Kostakos, "Huntgpt: Integrating machine learning-based anomaly detection and explainable ai with large language models (llms)," 2023. [Online]. Available: https://arxiv.org/abs/2309.16021

[3] J. Qi, S. Huang, Z. Luan, S. Yang, C. Fung, H. Yang, D. Qian, J. Shang, Z. Xiao, and Z. Wu, "Loggpt: Exploring chatgpt for log-based anomaly detection," in *2023 IEEE International Conference on High Performance Computing & Communications, Data Science & Systems, Smart City & Dependability in Sensor, Cloud & Big Data Systems & Application (HPCC/DSS/SmartCity/DependSys)*, 2023, pp. 273–280. [Online]. Available: https://doi.org/10.1109/HPCC-DSS-SmartCity-DependSys60770.2023.00045

[4] Y. Yao, J. Duan, K. Xu, Y. Cai, Z. Sun, and Y. Zhang, "A survey on large language model (llm) security and privacy: The good, the bad, and the ugly," *High-Confidence Computing*, vol. 4, no. 2, p. 100211, 2024. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2667295224000 14X

[5] K. Greshake, S. Abdelnabi, S. Mishra, C. Endres, T. Holz, and Fritz, "Not what you've signed up for: Compromising real-world llm-integrated applications with indirect prompt injection," in *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security*, ser. AISec '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 79–90. [Online]. Available: https://doi.org/10.1145/3605764.3623985

[6] A. Happe, A. Kaplan, and J. Cito, "Llms as hackers: Autonomous linux privilege escalation attacks," 2024. [Online]. Available: https://arxiv.org/abs/2310.11409

[7] L. Weidinger, J. Mellor, M. Rauh, C. Griffin, J. Uesato, P.-S. Huang, M. Cheng, M. Glaese, B. Balle, A. Kasirzadeh, Z. Kenton, S. Brown, W. Hawkins, T. Stepleton, C. Biles, A. Birhane, J. Haas, L. Rimell, L. A. Hendricks, W. Isaac, S. Legassick, G. Irving, and I. Gabriel, "Ethical and social risks of harm from language models," 2021. [Online]. Available: https://arxiv.org/abs/2112.04359

[8] E. Kasneci, K. Sessler, S. Küchemann, M. Bannert, D. Dementieva, F. Fischer, U. Gasser, G. Groh, S. Günnemann, E. Hüllermeier, S. Krusche, G. Kutyniok, T. Michaeli, C. Nerdel, J. Pfeffer, O. Poquet, M. Sailer, A. Schmidt, T. Seidel, M. Stadler, J. Weller, J. Kuhn, and G. Kasneci, "Chatgpt for good? on opportunities and challenges of large language models for education," *Learning and Individual Differences*, vol. 103, p. 102274, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1041608023000195

[9] J. Qadir, "Engineering education in the era of chatgpt: Promise and pitfalls of generative ai for education," in *2023 IEEE Global Engineering Education Conference (EDUCON)*, 2023, pp. 1–9. [Online]. Available: https://doi.org/10.1109/EDUCON54358.2023.10125121

[10] R. Michel-Villarreal, E. Vilalta-Perdomo, D. E. Salinas-Navarro, R. Thierry-Aguilera, and F. S. Gerardou, "Challenges and opportunities of generative ai for higher education as explained by chatgpt," *Education Sciences*, vol. 13, no. 9, 2023. [Online]. Available: https://www.mdpi.com/2227-7102/13/9/856

[11] D. Baidoo-anu and L. Owusu Ansah, "Education in the era of generative artificial intelligence (ai): Understanding the potential benefits of chatgpt in promoting teaching and learning," *Journal of AI*, vol. 7, no. 1, p. 52–62, 2023. [Online]. Available: https://doi.org/10.61969/jai.1337500

[12] Y. Wu, "Integrating generative ai in education: how chatgpt brings challenges for future learning and teaching," *Journal of Advanced Research in Education*, vol. 2, no. 4, pp. 6–10, 2023. [Online]. Available: https://doi.org/10.56397/JARE.2023.07.02

[13] R. Liu, C. Zenke, C. Liu, A. Holmes, P. Thornton, and D. J. Malan, "Teaching cs50 with ai: Leveraging generative artificial intelligence in computer science education," in *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*, ser. SIGCSE 2024. New York, NY, USA: Association for Computing Machinery, 2024, p. 750–756. [Online]. Available: https://doi.org/10.1145/3626252.3630938

[14] J. Woodrow, A. Malik, and C. Piech, "Ai teaches the art of elegant coding: Timely, fair, and helpful style feedback in a global course," in *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*, ser. SIGCSE 2024. New York, NY, USA: Association for Computing Machinery, 2024, p. 1442–1448. [Online]. Available: https://doi.org/10.1145/3626252.3630773

[15] D. Cambaz and X. Zhang, "Use of ai-driven code generation models in teaching and learning programming: a systematic literature review," in *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*, ser. SIGCSE 2024. New York, NY, USA: Association for Computing Machinery, 2024, p. 172–178. [Online]. Available: https://doi.org/10.1145/3626252.3630958

[16] V. D. Kirova, C. S. Ku, J. R. Laracy, and T. J. Marlowe, "Software engineering education must adapt and evolve for an llm environment," in *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*, ser. SIGCSE 2024. New York, NY, USA: Association for Computing Machinery, 2024, p. 666–672. [Online]. Available: https://doi.org/10.1145/3626252.3630927

[17] Y. Shen, X. Ai, A. G. Soosai Raj, R. J. Leo John, and M. Syamkumar, "Implications of chatgpt for data science education," in *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*, ser. SIGCSE 2024. New York, NY, USA: Association for Computing Machinery, 2024, p. 1230–1236. [Online]. Available: https://doi.org/10.1145/3626252.3630874

[18] P. Pusey, M. Gupta, S. Mittal, and M. Abdelsalam, "An analysis of prerequisites for artificial intelligence/machine learning-assisted malware analysis learning modules," in *Journal of The Colloquium for Information Systems Security Education*, vol. 11, no. 1, 2024, pp. 5–5. [Online]. Available: https://doi.org/10.53735/cisse.v11i1.177

[19] C. Ubah, P. Zaleppa, B. Taylor, and S. Kaza, "Evaluation of ai models to update cybersecurity curriculum," in *Journal of The Colloquium for Information Systems Security Education*, vol. 11, no. 1, 2024, pp. 8–8. [Online]. Available: https://doi.org/10.53735/cisse.v11i1.183

[20] J. Su and W. Yang, "Unlocking the power of chatgpt: A framework for applying generative ai in education," *ECNU Review of Education*, vol. 6, no. 3, pp. 355–366, 2023. [Online]. Available: https://doi.org/10.1177/20965311231168423

[21] C. K. Y. Chan, "A comprehensive ai policy education framework for university teaching and learning," *International journal of educational technology in higher education*, vol. 20, no. 1, p. 38, 2023. [Online]. Available: https://doi.org/10.1186/s41239-023-00408-3

[22] P. Prasad and A. Sane, "A self-regulated learning framework using generative ai and its application in cs educational intervention design," in *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*, ser. SIGCSE 2024. New York, NY, USA: Association for Computing Machinery, 2024, p. 1070–1076. [Online]. Available: https://doi.org/10.1145/3626252.3630828

[23] OpenAI, "Chatgpt and openai api," 2024, accessed: 2024-09-06. [Online]. Available: https://openai.com/chatgpt

[24] Microsoft, "Microsoft 365 copilot: Ai assistant for productivity," 2024, accessed: 2024-09-06. [Online]. Available: https://copilot.microsoft.com/

[25] Anthropic, "Claude: Next generation ai assistant by anthropic," 2024. [Online]. Available: https://claude.ai/

[26] Gamma Tech, Inc., "Gamma: Ai-powered presentations and documents," 2024. [Online]. Available: https://gamma.app/

[27] Tome, "How to create a great presentation with ai in tome," 2024. [Online]. Available: https://tome.app/blog/how-to-create-a-great-presentation-with-ai-in-tome

[28] Google, "Google colaboratory: Jupyter notebooks in the cloud," 2024. [Online]. Available: https://colab.google/

[29] A. Kohn and S. D. Blum, *Ungrading: Why rating students undermines learning (and what to do instead)*. West Virginia University Press, 2020.

[30] D. Kokotsaki, V. Menzies, and A. Wiggins, "Project-based learning: A review of the literature," *Improving schools*, vol. 19, no. 3, pp. 267–277, 2016. [Online]. Available: https://doi.org/10.1177/1365480216659733

[31] M. Erkens and D. Bodemer, "Improving collaborative learning: Guiding knowledge exchange through the provision of information about learning partners and learning contents," *Computers & Education*, vol. 128, pp. 452–472, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S036013151830280X

APPENDIX

CSEC-559/659-03: Generative AI in Cybersecurity 2235 TLP: White

5 - Checklist Strategies for Small Businesses

Use the strategies below for reducing digital footprints as a checklist to minimize your small business's digital footprint.

For further information and strategies, visit the following resource:
https://bizee.com/blog/how-to-protect-your-digital-footprint-as-a-small-business-owner

5.1 - Getting to Know Your Business's Digital World

Your small business's digital presence is a cohesive network, where your website, social media, and online tools unite to enhance your online visibility. Together, they create a supportive ecosystem that propels your business forward.

- **Your Website:** This is your online storefront, the first place new friends (aka customers) visit to see what you're all about. It's where you share your story, what you offer, and how to get in touch. Keeping it fresh and welcoming is key to making a great first impression.
- **Social Media:** These are the spots where you hang out and chat with customers, share updates, and build a community. It's like being at a never-ending networking event, but in your comfy clothes. These platforms help you stay connected, offer support, and really get to know your audience.
- **Cloud Services:** Imagine having a super-efficient back-office that's always up in the cloud, taking care of storage, teamwork, and all the nitty-gritty details of running your business smoothly. These tools are behind-the-scenes heroes, making sure everything runs like clockwork at an affordable cost.

Together, these elements create your business's digital footprint. It's like your business's mark on the online world, helping you operate smoothly, reach more people, and manage your data smartly. But, it's also important to keep an eye on your digital neighborhood to ensure everything is safe and presents your business in the best light.

5.2 - Reducing Digital Footprint for Your Digital World

By implementing these strategies, small businesses can significantly reduce their digital footprint, enhancing both their security and the trust of their customers. This approach not only

Fig. 1. Sample student outcome from Project #1. In this sample, the students produced a document that guided small businesses toward securing their online footprint. Students used generative AI to produce draft material, followed by peer review and human expertise to polish the generative material into a clear and useful document.

Fig. 2. Sample student outcome from Project #2. In this sample, the students built a functional keylogging client and server in the Rust programming language, which was a language that the students had no prior experience with. The students were not constrained to using unfamiliar languages but were directed to explore generative AI's ability to expand their capabilities in whatever way they saw fit.

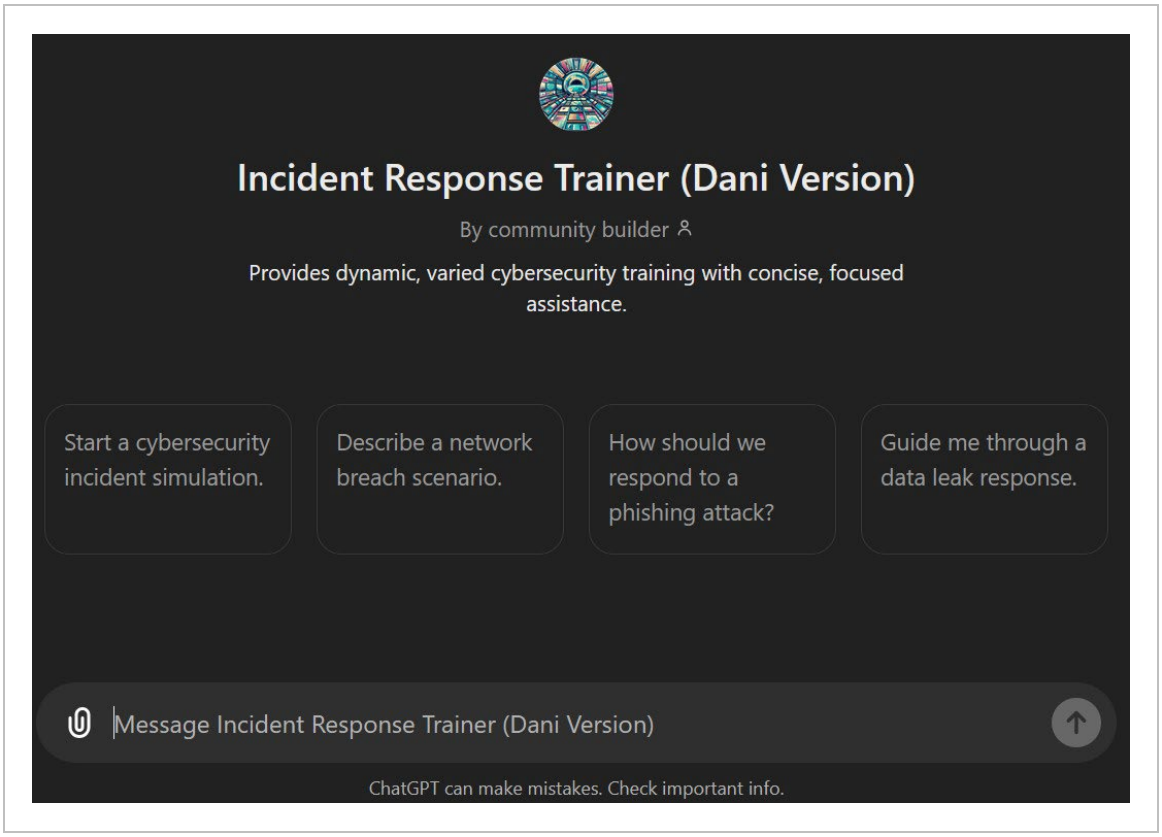

Fig. 3. Sample student outcome from Project #3. In this sample, the students built a customized chatbot using OpenAI's GPT builder for incident response training simulations. This customized chatbot included developing a hidden prompt to direct the bot and providing relevant supplemental materials to fill the bot's knowledge base.