Enhancing Cybersecurity Education with Artificial Intelligence Content

Fernando Brito
fbrit005@fiu.edu
Cyber-Physical Systems Security Lab,
School of Computing and Information Sciences,
Florida International University,
Miami, FL, USA

Yassine Mekdad
ymekdad@fiu.edu
Cyber-Physical Systems Security Lab,
School of Computing and Information Sciences,
Florida International University,
Miami, FL, USA

Monique Ross
ross.1982@osu.edu
Engineering Education Department,
The Ohio State University,
Columbus, OH, USA

Mark A. Finlayson
markaf@fiu.edu
Cognition, Narrative, and Culture Laboratory (Cognac Lab),
School of Computing and Information Sciences,
Florida International University,
Miami, FL, USA

Selcuk Uluagac
suluagac@fiu.edu
Cyber-Physical Systems Security Lab,
School of Computing and Information Sciences,
Florida International University,
Miami, FL, USA

Abstract

Artificial Intelligence (AI) has become a fundamental tool for cybersecurity researchers and practitioners. It is frequently used to address major security problems such as supply chain attacks, ransomware threats, and social engineering. In this context, integrating AI into cybersecurity workflows requires incorporating AI-driven approaches into the educational training of the cybersecurity workforce. This paradigm shift in academic settings will introduce the necessary skills for cybersecurity professionals to operate modern AI-based systems. Yet, the current cybersecurity curriculum still suffers from the absence of AI resources, particularly the detailed understanding of the appropriate AI mechanisms. Such absence leaves skill gaps for future professionals and practitioners in the industry. To address this, we designed an academic lecture module on AI covering both theory and practice. Then, we taught the module across six cybersecurity courses in our institution. To assess the effectiveness of integrating AI materials into cybersecurity education, we collected data by presenting two surveys before and after the lecture (concluding 81 participants per survey). Specifically, we utilized widely accepted models for unbiased analysis of our data. Our experimental results show positive AI knowledge improvement by 30% of the participants, demonstrating the beneficial impact of the lecture. Then, we observed a high similarity score between the survey responses and the lecture content, reaching 84%. Moreover, our sentiment analysis results reflect positive feedback from the participants with a positive score of 0.50. Overall, our study serves as a reference for instructional designers for developing educational curricula aiming to integrate AI into cybersecurity education.

SIGCSE TS 2025, February 26-March 1, 2025, Pittsburgh, PA, USA

ACM ISBN 979-8-4007-0531-1/25/02
https://doi.org/10.1145/3641554.3701958

CCS Concepts

• **Applied computing** → **Education**; • **Computing methodologies** → **Artificial intelligence**.

Keywords

Cybersecurity; Education; Artificial intelligence; Integration; Natural Language Processing; Machine Learning

ACM Reference Format:
Fernando Brito, Yassine Mekdad, Monique Ross, Mark A. Finlayson, and Selcuk Uluagac. 2025. Enhancing Cybersecurity Education with Artificial Intelligence Content. In *Proceedings of the 56th ACM Technical Symposium on Computer Science Education V. 1 (SIGCSE TS 2025), February 26-March 1, 2025, Pittsburgh, PA, USA*. ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3641554.3701958

1 Introduction

In the past few years, the landscape of cybersecurity education has gathered significant attention and widespread interest in response to the rising concerns of major security threats (e.g., ransomware attacks, phishing attacks, and Advanced Persistent Threats (APTs)) [25, 29, 35]. Despite considerable research efforts that have been focused on enabling a secure and trustworthy cyberspace using data-driven AI models (e.g. clustering, decision trees, support vector machines, artificial neural networks, and deep learning) [9, 17, 34], the curriculum guidelines established by several organizations such as ACM CCECC [31] and NIST NICE [21] have not yet embraced the integration of AI materials [3]. Therefore, incoming cybersecurity academics and practitioners might likely become unfamiliar with effective and novel techniques that are currently used by the cybersecurity workforce. Moreover, the current curriculum for cybersecurity education across academic institutions remains unchanged with its classical structure [3, 11]. Furthermore, studies

regarding cybersecurity education propose frameworks to optimize the curriculum by including new or commonly used approaches to address cyber threats (e.g., spam filtering, threat prediction, and threat identification) [6, 11].

More recently, cybersecurity practitioners and professionals have already incorporated AI into their security pipelines to enhance the protection of various systems (e.g., intrusion detection systems, endpoint security systems, and cloud infrastructure) [14, 18, 32]. Despite the frequent consideration of AI-based systems in the industry, only one study considers integrating AI into cybersecurity education [7]. It consists of lecturing basic AI concepts and testing the students on such materials. In this line of research, current educational approaches do not adequately demonstrate the application of AI-based techniques to the students. Therefore, we believe it is essential to effectively design AI educational materials for an inclusive and broader population, that could be swiftly integrated into an already packed cybersecurity curriculum.

In this paper, we are motivated to fill the skill gaps due to the absence of AI resources in cybersecurity education. As part of this effort, we designed an AI lecture module that systematically covers the relation of AI in cybersecurity in both theory and practice. The module covers several AI key concepts, including the necessary terminology, different types of learning, frequently used algorithms, and popular frameworks, most of which are commonly considered in the industry. Afterward, we conducted a usability study by teaching the AI lecture module across six cybersecurity graduate and undergraduate courses within our academic institution, involving 81 participants. We collected qualitative and quantitative data by presenting two surveys. We presented the first survey before the lecture while the second one directly after the lecture. Finally, we explored the potential impact of integrating AI into the cybersecurity curriculum by comprehensively evaluating our collected data. After conducting extensive data analysis using widely accepted Natural Language Processing (NLP) models, our experimental results show that the participants achieved an increase in AI knowledge by 30%. Such an increase underscores the participant's understanding of the topics and materials presented in the lecture. Further, we performed a topic modeling analysis by extracting keywords and identifying topics from textual data. Our results highlight the participants' concentration on understanding the provided AI lecture. Finally, we analyzed the participant's feedback and engagement. We performed sentiment analysis using a contextual model. Our experimental results show an average positive score of 0.50 from the RoBERTa model. Further analysis of feedback responses suggests potential improvements to the lecture. For example, several comments suggest lengthening the lecture duration or including more practical activities. This feedback verifies the participant's engagement and their suggestions about the lecture. Overall, our findings clearly demonstrate the feasibility of integrating AI materials into the cybersecurity curriculum.

Contributions: The main contributions of our work are as follows:

- We designed an AI lecture module that contains theoretical and practical AI material. The lecture's content includes AI key components such as the necessary terminology, common algorithms, and popular frameworks.
- We conducted a usability study by teaching the AI lecture module across six cybersecurity graduate and undergraduate courses within our academic institution, involving 81 participants.
- We demonstrate a 30% increase in AI knowledge among the participants. This outcome positively contributes to the cybersecurity curriculum. Moreover, we show through a topic distribution analysis the participant's concentration on understanding the provided lecture.
- By subsequently evaluating the participant's engagement and feedback, we show through a contextual model a positive average score of 0.50 using the RoBERTa model. This score signifies that the majority of the feedback is positive or contains positive sentiments.

Organization: The remainder of this paper is organized as follows: Section 2 provides a literature review on cybersecurity education and AI. In Section 3, we describe our research methodology. In Section 4, we provide our experimental results. Section 5 presents a discussion of our study. Finally, Section 6 concludes the paper.

2 Literature Review

Cybersecurity has become a subject of interest due to the development of cyber threats and their major concerns [5, 15, 23, 24, 26]. In response, specialists have developed novel detection and analysis strategies to analyze, prevent, and detect cyber threats [19, 22, 33]. To that end, several machine and deep learning models are used within the cybersecurity workforce for threat detection, analysis, monitoring, and prevention [12, 20, 28]. Despite its use by practitioners, AI has not been formally integrated into the cybersecurity education curriculum [3]. However, studies regarding cybersecurity education observe the effects of integrating AI into its courses. For example, Laato et al. [13] proposed two approaches for teaching AI in a cybersecurity course. The authors suggested mentioning AI only when it is relevant during the course. The authors recommended another approach that consists of teaching cybersecurity concepts from an AI-driven perspective. In another study, the authors assessed survey responses regarding at-risk security behavior in a business setting to develop an AI training framework [1]. The authors found that most participants were not aware that their actions caused possible security risks. In another work, Ansari et al. [2] proposed an AI-based cybersecurity training framework to prevent phishing attacks. In this work, the authors considered an established framework and a platform to create an AI-based cybersecurity awareness program. Such studies indicate that AI is applicable in cybersecurity training and education.

Difference from existing works: Differently from existing works, in this paper, we design an AI lecture module and perform a usability study across six graduate and undergraduate cybersecurity courses, involving 81 participants. In contrast to Laato et al. [13], which focuses on teaching cybersecurity from an AI perspective, our study collects qualitative and quantitative data to measure participants' AI knowledge improvement from a cybersecurity perspective. Unlike other methodologies, our methodology relies on presenting an in-depth AI lecture, pre-lecture, and post-lecture surveys to evaluate participants' performance. For instance, Farahmand et al. [7] does not examine participants' knowledge before providing the lecture.

However, we assess participants' AI knowledge before the lecture. Additionally, compared to Ansari [1], our methodology further differentiates itself by applying Natural Language Processing (NLP) techniques such as sentiment analysis, text extraction, and text preprocessing [27]. We use NLP to automate the survey evaluation process, eliminating bias from the student feedback and engagement analysis.

3 Methodology

In this section, we present our research methodology. Then, we explain our process for collecting the survey data. Further, we describe our text extraction, pre-lecture survey, post-lecture survey, and feedback analysis procedures.

3.1 Data Collection

For our data collection, we provided two surveys to students enrolled in cybersecurity courses within our academic institution. We presented two surveys within the lecture titled *"A Lecture on Artificial Intelligence, Machine Learning, and Deep Learning: From Theory to Practice"*. We gave the first survey to the students before providing the lecture. The first survey aims to assess the students' knowledge on AI and cybersecurity. It includes questions about demographics, cybersecurity, the basics of AI, models, and their corresponding performance metrics. After the students completed the first survey, we presented the lecture to the students. We organized the lecture into sections that contained material such as AI basics, foundations, popular frameworks, and types of learning including supervised and unsupervised models. The sections cover important information regarding AI, its origins, development, current models, and applications. Following the lecture, we provided the students with the second survey. The second survey is predominantly about AI metrics, AI models, Deep Learning, and AI training. The second survey measures the student's knowledge of AI after the lecture is provided. We designed both surveys to contain the same categories of questions. Below, we focus our assessments on three types of questions asked to the participants:

- **Open Response:** Participants are asked to answer an open-ended question.
- **Multiple Choice:** We tasked participants with choosing one answer among a group of set choices.
- **Select All That Apply:** We tasked respondents with selecting all answers that apply to them.

In the pre-lecture survey, we included questions about cybersecurity, computer science, and AI. In the cybersecurity questions, we assessed participants' knowledge of cybersecurity. Regarding the computer science questions, we evaluated participants' knowledge and experience with programming languages. Lastly, we used the AI questions to test if the participants had prior experience with AI development or usage. These questions also assess their familiarity with specific AI concepts (such as performance metrics and machine learning models).

In the post-lecture survey, we also included questions about cybersecurity, computer science, and AI. However, the survey largely consists of AI questions. The cybersecurity and computer science questions are related to AI topics such as performance metrics, AI models, and future use. The remainder of the questions are strictly about AI. These questions ask students to choose the correct AI model or techniques. Other open-ended questions assess their knowledge of performance metrics, AI frameworks, models, learning types, and applications. They also serve to analyze if they feel more comfortable with using AI in the future.

We presented both surveys using the Qualtrics platform. We distributed the surveys via anonymous links and QR codes. did not collect sensitive or personal data from the participants. We saved and exported the survey responses as CSV files for later analysis. The only responses that we considered were from students who completed both surveys. We sanitized our data while exporting the survey responses to a spreadsheet to mitigate missing or noisy data from our response collection. For example, we corrected incorrect characters and filled the empty fields with "N/A". This process would limit the threats to validity, and guarantee the reliability of our collected data. In total, 81 participants completed both surveys. We found that in our demographic distribution, 77.8% of participants are male, 19.8% are female, and 2% identified as other or did not specify. Additionally, 43.2% of participants are Hispanic and 49.4% are white.

IRB Approval: We submitted our methodology and survey questions to our university's Institutional Review Board. Our work is exempted under category #1 by the Human Subject Research Protection department, as we did not collect sensitive or personally identifiable information from the participants. The participants were shown a consent letter to inform them about the purpose of our research and their tasks if they chose to participate.

3.2 Text Extraction and Survey Analysis

After we collected our survey data, we began our data analysis and text extraction. We extracted the lecture text to compare our collected survey responses with the lecture. In order to do this, we organized our methodology into two stages as illustrated in Figure 1. These two stages are the *Lecture Extraction* and *Survey Analysis* stages. Only the *Survey Analysis* stage requires the survey responses from the data collection. Moreover, our *Feedback Analysis* is independent of the other two stages and can be considered separately.

First, we extracted the text from the lecture in the *Lecture Extraction* stage. The goal of this stage is to extract the lecture for future comparison with the survey responses. In this stage, we converted the lecture slides into a text file. Since the lecture materials are in the form of a PDF, we extracted the text through the library PyPDF2 [30]. This library contains functions for extracting and processing text from PDF files. Once we extracted the text from the lecture materials, we used a Natural Language Toolkit (NLTK) [4] to tokenize each word in the text. Finally, we appended the list of tokenized words together to finalize the text as a string and write it into a text file. Following the *Lecture Extraction* stage, we continue into the *Survey Analysis* stage. In this stage, we extracted topics from the survey responses and evaluated the similarity of the responses to the lecture. We obtained the necessary results from this survey analysis. However, it is necessary to clean the data beforehand. We removed unnecessary rows and columns from the CSV files. The removed rows and columns contain non-critical data provided by the Qualtrics platform (e.g., date of completion, time until

completion, and geographical coordinates). After the preliminary data cleaning, we prepared the CSV file with the correct answers. We scored survey responses based on similarity with the CSV. We use the spaCy library to compare the similarity of the two texts [10]. Then, we scrutinized the data and removed outliers that contained unnecessary data. We removed these outliers as they can negatively impact the data and provide inaccurate results. We performed this process in both the pre-lecture and post-lecture survey responses. In the *Survey Analysis* stage, we compared the survey responses to the extracted lecture text. Similarly to the previous step, we calculated the average similarity scores of both surveys to the extracted lecture text. Then, we subtracted the averages to verify variations in similarity after the lecture. In this stage, we also performed a topic extraction. We extracted keywords from the survey responses and sorted them into topics. We used the BERTopic model to extract these topics. This model uses a pre-trained language model to extract and distribute topics into groups [8]. To mitigate potential errors during the NLP analysis, we implemented several preprocessing steps, notably removing the stopwords and correcting incorrect characters in the data. Furthermore, we considered a popular sentence transformer model "all-MiniLM-L6-v2" within the BERTopic. It enables effective topic modeling given its robustness in terms of capturing contextual information. Using the BERTopic model, we visualized and ranked the extracted topics by size. Then, we compared the extracted topics to the survey-to-lecture similarity scores to verify those results.

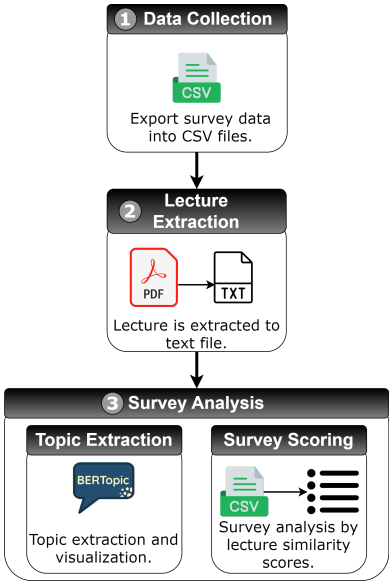

Figure 1: Diagram of our methodology. The diagram shows the relationship between *Lecture Extraction* and *Survey Analysis* and displays the remaining stage. The *Feedback Analysis* is done independently from the other two stages.

3.3 Feedback Analysis

In the *Feedback Analysis*, we analyzed the student feedback and engagement using the post-lecture survey questions. To analyze the feedback we performed a sentiment analysis using the RoBERTa model. Although the considered NLP-based implementation in our methodology provides a bias-free analysis, it does have some limitations for analyzing open-ended survey responses. For example, numerical ratings or null values in the data might not be properly interpreted. However, these limitations are commonly accepted by the majority of the studies using NLP techniques. RoBERTa is a pre-trained model based on the BERT model [16]. It can output positive, neutral, and negative scores. It analyzes the sentimentality of a text while considering its context. We calculated the average sentimentality scores from the participants' feedback. Then, we used the average scores to find the general sentimentality regarding the feedback.

4 Experimental Results

In this section, we provide our quantitative and qualitative results from the *Survey Analysis* and *Feedback Analysis* stages. These results are verified through the use of several language models or by similarity analysis.

4.1 Topic Distribution Results

In the pre-lecture survey, we included a few cybersecurity-based questions to measure the participants' knowledge about cybersecurity. Our experimental results show that, on average, the pre-lecture cybersecurity assessment percentage score is 84.39%. These results clearly confirm that the students are familiar with cybersecurity concepts. Therefore, we acknowledge that the participants are well-informed about cybersecurity concepts and are indeed cybersecurity students. Based on topic extraction observations, we find that the largest topics (by size) are related to survey response keywords or AI. The largest topics in the pre-lecture and post-lecture distributions contain 96 and 51 keywords respectively. In Table 1, we present the top ten largest topics from the pre-lecture and post-lecture results. Through the BERTopic model, we extract the most common keywords in the survey responses and represent them as topics. The topic distributions are organized as clusters based on their similarity to other topics. The pre-lecture survey clusters are related to programming, cybersecurity, chatbots, AI basics, and NLP. Other topic clusters are formed due to uncategorizable keywords in the survey responses of participants. These topic clusters can be considered as miscellaneous topics. In contrast, the post-lecture topic clusters have higher specificity than the pre-lecture clusters. In the post-lecture topic extraction, we find that the largest topic is related to decision tree algorithms (keywords "random", "forest", "tree", "decision", and "trees"). This topic is more specific in comparison to the pre-lecture topic extraction (keywords "can", "it", "tasks", "to", "help" and "ai", "ml", "and", "in", "to"). Additionally, the distributions shown in the post-lecture extraction, in Table 1, are related to AI models, performance metrics, AI applications, Deep Learning, cybersecurity, neural networks, and AI frameworks. The post-lecture topic distribution results show more specificity and variation in comparison to the pre-lecture results.

4.2 Survey Analysis Results

To analyze the extracted surveys, we compared the survey responses to the CSV file with correct answers and the extracted

Pre-lecture Survey Topics	Number of Keywords
can | it | tasks | to | help	96
ai | ml | and | in | to	94
user | identity | verifying | process | authentication	82
no | nope | yes | marked | exp	72
concept | this | know | do | not	54
java | python | languages | programming | my	48
chatbots | text | speech | translation | sentiment	41
do | know | not	36
chatgpt | gpt | chat | besides | bixby	35
comfortable | feel | very | would | not	35
Post-lecture Survey Topics	
random | forest | tree | decision | trees	51
sure | not | familiar | am | entirely	48
ai | ml | of | be | to	44
no | practically | its | really | good	43
unsupervised | supervised | prediction | sets | semi	42
points | clusters | cluster | centroids | classified	36
text | translation | sentiment | language | analysis	36
reinforcement | supervised | learning | unsupervised | semi	34
yes | wonderful | access | now | drive	34
accuracy | model | bias | improve | accurate	32

Table 1: Top 10 pre-lecture and post-lecture topics. These topics are groups of keywords that are used similarly within a text. This most often occurs when keywords are used together frequently.

text from the lecture. As a result, our survey analysis has two similarity scores, The first similarity score is the survey-to-survey similarity. This method of analysis is used to grade the post-lecture survey results similarly to an exam. The second method of analysis is the survey-to-lecture analysis. This analysis measures student improvement in AI knowledge in the post-lecture survey.

In the survey-to-survey similarity scores, 71 students achieved passing grades. We consider a grade to be passing if they could score 70% or higher in the post-lecture survey-to-survey analysis. Further observations show that 3 students scored between 70-80%, 11 students ranged between 80-90%, and 57 students ranged between 80-90%. As a result, most students passed the survey analysis. Unfortunately, the remainder of the students did not achieve a score greater than or equal to 70% or did not answer the post-lecture survey correctly (e.g., "N/A", "I don't know", etc.). Finally, the survey-to-lecture similarity scores show a positive trend toward the post-lecture scores. The pre-lecture results' average similarity score to the AI lecture was 79.30%. The second survey has an average similarity score of 84.44%, an increase of 5.14%. Furthermore, we find that the maximum performance increase is 30% while the lowest is 2%. However, it is important to note that this average score differs from the cybersecurity assessment average as they do not evaluate the same topics. The increase in the survey-to-lecture similarity scores in the post-lecture surveys shows improvement in the participants' knowledge of AI after the lecture. It also indicates a positive trend, such that the participants retain knowledge of the AI lecture.

4.3 Feedback Analysis Results

In the *Feedback Analysis*, we find that the consensus from the participants' feedback and engagement has positive sentimentality. As illustrated in Figure 2, the RoBERTa model shows that the highest average score in its sentimentality analysis is the average positive score of 0.50. However, while Figure 2 shows that the positive score is much higher than the negative score, the neutral score, 0.45, is close to the positive score.

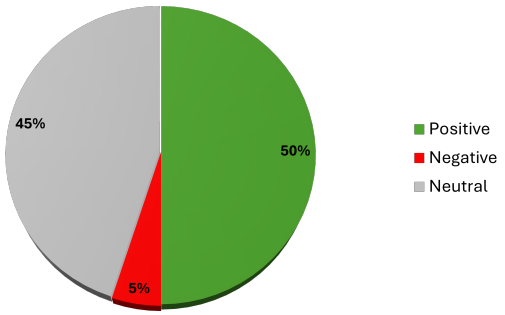

Figure 2: Pie chart illustrating average percentage scores from RoBERTa sentiment analysis.

In the RoBERTa analysis, the highest score is the positive average, while the second highest is the neutral average. This could be a result of comments regarding the lecture. For instance, one participant said:

Participant #71: *"I would increase the lecture duration to touch on additional concepts. Although outside of the scope of a lecture, providing students with the tools and exercises to implement machine learning using test datasets, or giving a task where they are told to extract features or create a folder that will be used for DL auto-feature extraction, would provide more practical experience with these tools and give a better perspective of their performance and use cases".*

Another participant stated:

Participant #9: *"I would make it maybe more interactive with the students, just because this study might be new to some people and they might not fully grasp the concepts from the beginning. Make sure to keep them engaged".*

The remarks shown are not specifically negative. In other words, they are using language that is not viewed as negative. However, they do present some dissatisfaction with the lecture length or engagement. Consequently, these comments may be considered neutral (neither positive nor negative) by the sentiment analysis model. Thus, such remarks could be the source of the high neutrality score in the RoBERTa analysis. Based on the feedback comments, negative feedback is not aimed towards the lecturer. However, the feedback does concern the lecture. A sample of comments reflects their displeasure with aspects of the lecture. For example, one participant states:

Participant #1: *"It's hard to learn about this in such a short time".*

This particular participant expressed how the lecture was too short for the quantity of AI material covered. Likewise, other participants shared similar feelings and would have preferred that the lecture be longer or broader. Such feedback may be viewed as negative or neutral depending on prior context by the RoBERTa model. Regardless, the feedback comments attributed to the neutral and negative scores in the feedback analysis. Furthermore, the feedback towards the lecturer also displays some dissatisfaction. In the post-lecture survey, we asked participants if the lecturer performed well. Most participants acknowledge a high satisfaction regarding the presentation. Despite this, some comments detailed their suggestions towards the lecturer. One participant suggests:

Participant #35: *"I would interact a little bit more with the audience to gauge their understanding of the [foundational] concepts of computer science before engaging in such lecture. Some students might have a high CS background while others might have a focus on an applied branch and might not be very familiar with the math and the vocabulary for this lecture".*

The participant's feedback shares similarities to other comments. Several comments expressed that the lecturer should present more interactively with the participants. Others suggest that the lecturer should apply more hands-on techniques. Such responses may also increase the neutral scores in the sentiment analysis or appear as negative.

5 Discussion

According to our literature review, the provided studies only present minimal AI-centric content with limited depth and specificity [1, 7, 13]. Additionally, their suggested methodologies ignore testing the participants' knowledge before exposing them to the lecture. In contrast, we provide the students with an AI lecture that delves into the history, applications, and models of AI. Likewise, our surveys are designed to test students' familiarity with AI before and after the lecture. They also serve to record changes in performance regarding AI knowledge after the lecture. Our survey analysis results indicate an expected improvement in the survey scores after the AI lecture. While the total average performance increase in the survey results is 5.14%, we found that the largest performance increase after the lecture can be as high as 30%. A large performance increase indicates a possibility that a longer lecture duration or quantity of material shown could potentially further increase the student's familiarity and knowledge of AI. This change can be implemented as a longer lecture or several lectures presented throughout the scholastic semester. These modifications to the lecture can be applied in future studies. In fact, program directors and instructors at other institutions can customize and expand our AI-infused module over large elements of their curriculum. For instance, foundational AI concepts from our module, such as machine learning algorithms and its corresponding terminology can be swiftly integrated into the existing cybersecurity courses.

Additionally, we found fewer miscellaneous topics in the post-lecture topic distribution than in the pre-lecture results. The pre-lecture topic distribution shows more topics with keywords that are uncategorizable (e.g., "can", "it", "and", "no", "nope", "concept", "this", "do", and "know"). However, in the post-lecture results, there are fewer miscellaneous topics and clusters across the distribution. The higher specificity and absence of miscellaneous topics in the post-lecture results are proof of high confidence in the participants' responses. Lastly, the feedback and student engagement analysis presents student suggestions that may be used to improve the lecture. While several participants had positive comments about the lecture, others commented on possible suggestions to consider. These comments recommend that the lecture duration be extended or that the content be simplified since our lecture duration was one hour in length. However, some participant feedback responses explain that one hour is too short for the density and depth of the content within the lecture. As a result, the participants suggest lengthening the lecture or broadening the topics. Other post-lecture feedback responses advise the lecturer to proactively interact with the students. They also suggest that the lecturer should provide more practical examples to students. These remarks show that the lecturer could improve the quality of the presentation by asking the students questions or interacting with them more often. The participants also suggested that they should be given more practical demonstrations and examples to enhance the quality of their learning experience. We find this feedback valuable as it will increase the quality and reproducibility of future lectures.

6 Conclusion

In this paper, we designed an academic AI lecture module that includes both theoretical and practical content. We implemented this lecture module across six cybersecurity courses within our academic institution, involving 81 participants. Then, we conducted a usability study and assessed changes in students' performance after completing the module. Our results demonstrate a 30% increase in AI knowledge among participants. Additionally, we found that the topics discussed in the post-lecture survey were more diverse and specific compared to the pre-lecture survey, demonstrating enhanced understanding. The feedback and student engagement have also demonstrated positive sentiment, with the RoBERTa model showing an average positive score of 0.50. The negative sentiments detected by the model primarily pertained to suggestions for improving the lecture's content and delivery. In future studies, we suggest addressing these feedback comments by either lengthening the lecture duration or presenting the module throughout an academic semester to increase student performance. Lastly, due to the positive trend in survey-to-lecture and feedback sentimentality scores, the lecture successfully presented students with AI topics and techniques. Therefore, we conclude that it is feasible and beneficial to implement the AI module into cybersecurity education.

7 Acknowledgement

We thank the anonymous reviewers for their helpful feedback and time. This work was partially supported via US National Science Foundation's Intergovernmental Personnel Act Independent Research & Development Program, NSF Award DGE-2039606, US National Security Agency Award No. H982302110324, and Cyber Florida. The views expressed are those of the authors only, not of the funding agencies.

References

[1] Meraj Farheen Ansari. 2022. A quantitative study of risk scores and the effectiveness of AI-based Cybersecurity Awareness Training Programs. *International Journal of Smart Sensor and Adhoc Network* 3, 3 (2022), 1.

[2] Meraj Farheen Ansari, Pawan Kumar Sharma, and Bibhu Dash. 2022. Prevention of phishing attacks using AI-based Cybersecurity Awareness Training. *Prevention* 3, 6 (2022).

[3] AHMET ARIS, Luis Puche Rondon, Daniel Ortiz, Monique Ross, and Mark Finlayson. 2022. Integrating Artificial Intelligence into Cybersecurity Curriculum: New Perspectives. In *2022 ASEE Annual Conference & Exposition*.

[4] Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural language processing with Python: analyzing text with the natural language toolkit.* " O'Reilly Media, Inc.".

[5] Derin Cayir, Abbas Acar, Riccardo Lazzeretti, Marco Angelini, Mauro Conti, and Selcuk Uluagac. 2024. Augmenting Security and Privacy in the Virtual Realm: An Analysis of Extended Reality Devices. *IEEE Security & Privacy* 22, 1 (2024), 10–23.

[6] James Crabb, Christopher Hundhausen, and Assefaw Gebremedhin. 2024. A Critical Review of Cybersecurity Education in the United States. In *Proceedings of the 55th ACM Technical Symposium on Computer Science Education V. 1*. 241–247.

[7] Fariborz Farahmand. 2021. Integrating cybersecurity and artificial intelligence research in engineering and computer science education. *IEEE Security & Privacy* 19, 6 (2021), 104–110.

[8] Maarten Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. *arXiv preprint arXiv:2203.05794* (2022).

[9] Shuchi Grover, Brian Broll, and Derek Babb. 2023. Cybersecurity education in the age of ai: Integrating ai learning into cybersecurity high school curricula. In *Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1*. 980–986.

[10] Matthew Honnibal and Ines Montani. 2016. spaCy. https://github.com/explosion/spaCy.

[11] Eunyoung Kim and Razvan Beuran. 2018. On designing a cybersecurity educational program for higher education. In *Proceedings of the 10th International Conference on Education Technology and Computers*. 195–200.

[12] Murat Kuzlu, Corinne Fair, and Ozgur Guler. 2021. Role of artificial intelligence in the Internet of Things (IoT) cybersecurity. *Discover Internet of things* 1, 1 (2021), 7.

[13] Samuli Laato, Ali Farooq, Henri Tenhunen, Tinja Pitkamaki, Antti Hakkala, and Antti Airola. 2020. AI in Cybersecurity Education- A Systematic Literature Review of Studies on Cybersecurity MOOCs. (2020), 6–10.

[14] Jian-hua Li. 2018. Cyber security meets artificial intelligence: a survey. *Frontiers of Information Technology & Electronic Engineering* 19, 12 (2018), 1462–1474.

[15] Yuchong Li and Qinghui Liu. 2021. A comprehensive review study of cyber-attacks and cyber security; Emerging trends and recent developments. *Energy Reports* 7 (2021), 8176–8186.

[16] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).

[17] Javier Martínez Torres, Carla Iglesias Comesaña, and Paulino J García-Nieto. 2019. Machine learning techniques applied to cybersecurity. *International Journal of Machine Learning and Cybernetics* 10, 10 (2019), 2823–2836.

[18] Yassine Mekdad, Giuseppe Bernieri, Mauro Conti, and Abdeslam El Fergougui. 2021. The rise of ICS malware: A comparative analysis. In *European Symposium on Research in Computer Security*. Springer International Publishing, Cham, 496–511.

[19] Yassine Mekdad, Giuseppe Bernieri, Mauro Conti, and Abdeslam El Fergougui. 2021. A threat model method for ICS malware: the TRISIS case. In *Proceedings of the 18th ACM International Conference on Computing Frontiers*. Association for Computing Machinery, New York, NY, USA, 221–228.

[20] Yassine Mekdad, Faraz Naseem, Ahmet Aris, Harun Oz, Abbas Acar, Leonardo Babun, Selcuk Uluagac, Güliz Seray Tuncay, and Nasir Ghani. 2024. On the Robustness of Image-Based Malware Detection Against Adversarial Attacks. In *Network Security Empowered by Artificial Intelligence*. Springer Nature Switzerland, Cham, 355–375.

[21] William Newhouse, Stephanie Keith, Benjamin Scribner, and Greg Witte. 2017. National initiative for cybersecurity education (NICE) cybersecurity workforce framework. *NIST special publication* 800, 2017 (2017), 181.

[22] Ehsan Nowroozi, Mohammadreza Mohammadi, Pargol Golmohammadi, Yassine Mekdad, Mauro Conti, and A Selcuk Uluagac. 2023. Resisting deep learning models against adversarial attack transferability via feature randomization. *IEEE Transactions on Services Computing* 17, 1 (2023), 18–29.

[23] Ehsan Nowroozi, Seyedsadra Seyedshoari, Yassine Mekdad, Erkay Savaş, and Mauro Conti. 2022. Cryptocurrency wallets: assessment and security. In *Blockchain for Cybersecurity in Cyber-Physical Systems*. Springer International Publishing, Cham, 1–19.

[24] Harun Oz, Abbas Acar, Ahmet Aris, Güliz Seray Tuncay, Amin Kharraz, and Selcuk Uluagac. 2024. (In) Security of File Uploads in Node. js. In *Proceedings of the ACM on Web Conference 2024*. 1573–1584.

[25] Harun Oz, Ahmet Aris, Abbas Acar, Güliz Seray Tuncay, Leonardo Babun, and Selcuk Uluagac. 2023. {RøB}: Ransomware over Modern Web Browsers. In *32nd USENIX Security Symposium (USENIX Security 23)*. 7073–7090.

[26] Harun Oz, Daniele Cono D'Elia, Güliz Seray Tuncay, Abbas Acar, Riccardo Lazzeretti, and Selcuk Uluagac. 2024. With Great Power Comes Great Responsibility: Security and Privacy Issues of Modern Browser Application Programming Interfaces. *IEEE Security & Privacy* (2024).

[27] Rajvardhan Patil, Sorio Boit, Venkat Gudivada, and Jagadeesh Nandigam. 2023. A survey of text representation and embedding techniques in nlp. *IEEE Access* 11 (2023), 36120–36146.

[28] Iqbal H Sarker, Md Hasan Furhad, and Raza Nowrozy. 2021. Ai-driven cybersecurity: an overview, security intelligence modeling and research directions. *SN Computer Science* 2, 3 (2021), 173.

[29] Jayesh Soni, Surya Sirigineedi, Krishna Sai Vutukuru, SS ChandanaEswari Sirigineedi, Nagarajan Prabakar, and Himanshu Upadhyay. 2023. Learning-Based Model for Phishing Attack Detection. In *Artificial Intelligence in Cyber Security: Theories and Applications*. Springer, 113–124.

[30] Matthew Stamy. 2012. PyPDF2: https://github.com/mstamy2/PyPDF2.

[31] Cara Tang, Cindy Tucker, Christian Servin, and Markus Geissler. 2018. Computer science curricular guidance for associate-degree transfer programs. In *Proceedings of the 49th ACM Technical Symposium on Computer Science Education*. 435–440.

[32] Feng Tao, Muhammad Shoaib Akhtar, and Zhang Jiayuan. 2021. The future of artificial intelligence in cybersecurity: A comprehensive survey. *EAI Endorsed Transactions on Creative Technologies* 8, 28 (2021), e3–e3.

[33] Chandra Sekar Veerappan, Peter Loh Kok Keong, Zhaohui Tang, and Forest Tan. 2018. Taxonomy on malware evasion countermeasures techniques. In *2018 IEEE 4th World Forum on Internet of Things (WF-IoT)*. IEEE, 558–563.

[34] Vivek Verma, Krishna Sai Vutukuru, Sai Srinivas Divvela, and Surya Srikar Sirigineedi. 2022. Internet of things and machine learning application for a remotely operated wetland siphon system during hurricanes. In *Water Resources Management and Sustainability*. Springer, 443–462.

[35] Ibrar Yaqoob, Ejaz Ahmed, Muhammad Habib ur Rehman, Abdelmuttlib Ibrahim Abdalla Ahmed, Mohammed Ali Al-garadi, Muhammad Imran, and Mohsen Guizani. 2017. The rise of ransomware and emerging security challenges in the Internet of Things. *Computer Networks* 129 (2017), 444–458.