Toward Experiential Security and Privacy Training for AI Practitioners

Mujtaba Nazari
Department of Computer Science
Loyola University Chicago
Chicago, IL, USA
mnazari@luc.edu

Eric Chan-Tin
Department of Computer Science
Loyola University Chicago
Chicago, IL, USA
dchantin@luc.edu

Loretta Stalans
Department of Criminal Justice & Criminology
Loyola University Chicago
Chicago, IL, USA
Lstalan@luc.edu

Mohammed Abuhamad
Department of Computer Science
Loyola University Chicago
Chicago, IL, USA
mabuhamad@luc.edu

***Abstract*—The rapid adoption of artificial intelligence across industries has outpaced security and privacy training for AI practitioners. This paper presents methods, modules, and findings from an experiential training program designed to address security and privacy challenges in AI systems development and deployment. We conducted two program iterations: a comprehensive 12-workshop series (May-October 2024) and a condensed 6-workshop format (January-February 2025). The program combined expert-led panel sessions with hands-on laboratory activities, engaging 78 participants from diverse professional backgrounds. Evaluation through pre- and post-evaluation surveys and qualitative observations revealed improvements in cybersecurity knowledge and AI security awareness. Participants demonstrated enhanced ability to identify vulnerabilities, implement security measures, and develop organizational policies for AI-related risk mitigation. The condensed format showed comparable learning outcomes with improved completion rates. This effort highlights the increased need to establish cybersecurity and privacy training for AI professionals to develop secure and trustworthy AI systems.**

***Index Terms*—AI security, cybersecurity training, privacy-preserving AI, experiential learning, professional development, adversarial machine learning**

I. Introduction

Artificial intelligence systems have become integral to modern business operations and critical infrastructure. Organizations increasingly deploy machine learning models, generative AI tools, and automated decision systems without adequate understanding of associated security and privacy risks. Even AI developers with strong technical skills often lack formal training in cybersecurity principles and practices. This gap between AI development/adoption and security awareness creates vulnerabilities that adversaries actively exploit through adversarial attacks, such as data poisoning, model extraction, and privacy breaches [1]. Current cybersecurity training programs focus primarily on traditional network and system security, leaving AI-specific threats inadequately addressed. Academic curricula often treat AI security as an advanced topic, while industry professionals require immediate, practical knowledge to secure existing AI pipelines and deployments. The shortage of qualified AI security practitioners further exacerbates this skills gap, as demand for AI security expertise continues to outpace supply.

Besides these challenges, other challenges arise in the context of professional education as traditional online courses and certification programs often rely on passive learning and/or rigid schedules. Furthermore, they provide insufficient hands-on experience and lacking the flexibility and real-world relevance that AI professionals need to master complex concepts in security and privacy.

This paper presents the design and implementation of an experiential training program specifically designed for AI practitioners. The program addresses the identified gaps through expert-led panels, practical laboratory sessions, and real-world case studies. We report results from two program iterations, comparing effectiveness between comprehensive and condensed formats. Our evaluation methodology includes quantitative assessment of knowledge gains and qualitative analysis of skill development and organizational impact. The primary contributions of this work include: ➊ design and implementation of a comprehensive AI security training program, ➋ evaluation of experiential learning effectiveness for cybersecurity education, ➌ comparison of different program delivery formats, and ➍ recommendations and insights for scaling professional AI security training programs.

II. Related Work

The intersection of cybersecurity education and artificial intelligence represents a critical challenges facing modern educational institutions and organizations. Traditional cybersecurity awareness and training programs are foundational to organizational security postures. However, they are increasingly recognized as insufficient. They can not address the sophisticated threat landscape of this AI era [2]. The rapid adoption of AI technologies in workplace and educational settings. This has introduced novel security and privacy challenges. Existing training frameworks are not equipped to address these challenges [3]. Prior research related to this work spans three key areas: empirical studies examining the efficacy of conventional cybersecurity awareness programs, research on leveraging AI technologies to improve training outcomes, and studies addressing new risks introduced by AI adoption in educational and workplace contexts.

Recent research by [2] demonstrates that traditional cybersecurity training is ineffective. Their eight-month longitudinal study found no correlation between recent training completion and employees' ability to avoid phishing attacks. This challenges fundamental assumptions about standard training approaches. Meta-analytical research [4] shows traditional cybersecurity training has positive effects. However, these effects are insufficient. Critical limitations exist. These include lack of engagement, static content delivery, insufficient behavioral focus, and limited knowledge retention.

Research by [5] demonstrates how Large Language Models can transform cybersecurity training. They achieved this through personalization. Their study using GPT models showed significant improvements. These included tailored learning experiences, job-specific scenarios, dynamic content adjustment based on risk scores, and accurate delivery of technical concepts with policy references. Research by [6] demonstrated the effectiveness of AI-powered adaptive cybersecurity training in industrial environments. Their simulation-based study with 100 industrial employees showed a 72% reduction in phishing susceptibility. It also demonstrated 50% improvement in incident response time, and 69% increase in threat detection accuracy.

The adoption of AI tools in workplace environments has created unprecedented security and privacy challenges. Research by [3] identifies critical gaps in AI risk preparedness. The study shows 35% of security breaches involve employee activities, widespread use of unvetted LLMs. Traditional training programs lacking coverage of AI-specific threats like prompt injection and data exfiltration. Research by [7] reveals critical vulnerabilities in educational institutions. Over 1,600 school district data breaches since 2016. Education ranking third for hackers. Teachers widely used unauthorized AI tool. Additionally 43% of parents remain unaware of AI integration despite 46% of high school students using AI daily.

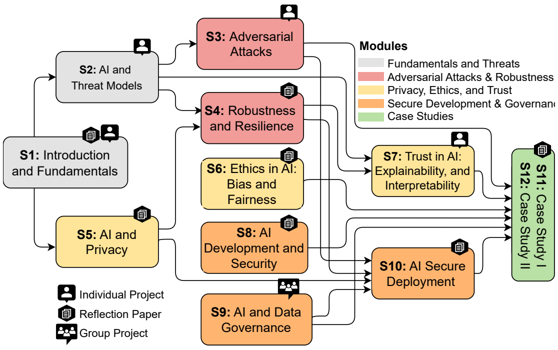
S1: Introduction and Fundamentals
S2: AI and Threat Models
S3: Adversarial Attacks
S4: Robustness and Resilience
S5: AI and Privacy
S6: Ethics in AI: Bias and Fairness
S7: Trust in AI: Explainability, and Interpretability
S8: AI Development and Security
S9: AI and Data Governance
S10: AI Secure Deployment
S11: Case Study I
S12: Case Study II
Modules
Fundamentals and Threats
Adversarial Attacks & Robustness
Privacy, Ethics, and Trust
Secure Development & Governan
Case Studies
Individual Project
Reflection Paper
Group Project

Fig. 1: The figure shows program workshops in different colors that represent their modules. To achieve optimal learning outcomes, participants are encouraged to follow the dependency graph and conduct the required lab activities.

III. SecureAI: Methods

A. Program Design

We propose a program that follows an experiential learning model. It combines theoretical knowledge with practical application. Each workshop consists of two components: expert-led panel discussions and hands-on laboratory activities. This structure ensures participants gain both conceptual understanding and practical skills. The modules and workshops are shown in Figure 1. The labs and hands-on activities are summarized in Table I.

Module 1: Fundamentals and Threats. includes Workshop 1 (Introduction and Fundamentals) and Workshop 2 (AI and Threat Models). These sessions establish foundational knowledge of AI systems and security threat models.

Module 2: Adversarial Attacks and Robustness. encompasses Workshop 3 (Adversarial Attacks) and Workshop 4 (Robustness and Resilience). Participants learn about attack vectors and defenses.

Module 3: Privacy, Ethics, and Trust. includes Workshop 5 (AI and Privacy: Differential Privacy and Federated Learning), Workshop 6 (Ethics in AI: Bias and Fairness), and Workshop 7 (Trust in AI: Transparency, Explainability, and Interpretability). These sessions address privacy-preserving techniques, ethical considerations, and building trustworthy AI systems.

Module 4: Secure Development and Data Governance. includes Workshop 8 (AI Development and Security) and Workshop 9 (AI and Data Governance: Regulations and Standards), and Workshop 10 (Secure Deployment and Operation of AI Systems). These sessions cover secure coding practices, compliance with data protection regulations, and strategies for deploying AI systems securely.

TABLE I: Lab sessions accompanied by workshops

Session	Key Activities
Lab 1	Identify cybersecurity risks in AI systems Assess risk likelihood and impact Propose mitigation strategies Implement security measures
Lab 2	Define white-box threat model Implement attacks (PGD, FGSM, C&W) Evaluate attack effectiveness Implement defenses (input sanitization, adversarial training)
Lab 3	Define black-box threat model Implement attacks (SimBA, MGAattack) Implement defenses Deploy adversarial detectors
Lab 4	Implement evaluation metrics Apply denoising methods Implement domain adaptation/transfer learning Assess robustness
Lab 5	Set up federated learning environment Implement client-side training Configure model aggregation Evaluate performance and privacy
Lab 6	Identify bias in datasets Measure bias using metrics Apply bias mitigation techniques Analyze impact on model performance
Lab 7	Analyze feature importance Implement interpretable models Apply interpretation methods (*e.g.,* GRAD & CAM)
Lab 8	Set up secure development environment Implement secure coding practices Deploy models securely Establish monitoring procedures
Lab 9	Assess data handling practices Develop data governance framework Apply compliance checklists Review industry standards
Lab 10	Containerize AI models Apply secure API design Implement access controls
Lab 11 & 12	Analyze real-world case studies Review security vulnerabilities Implement mitigation strategies Document lessons learned Final reflection

Module 5: Case Studies. includes Workshop 11 (Case Study #1) and Workshop 12 (Case Study #2), applying learned concepts to real-world scenarios. Participants analyze case studies to identify security vulnerabilities and propose mitigation strategies.

B. Program Implementation

We implemented the program in two distinct runs. The first run adopted the comprehensive 12-workshop format. The second run condensed the curriculum into 6 workshops over six weeks. Both formats maintained the same core content and learning objectives. Each workshop combined a 60-minute expert panel/lecture with a 60-minute hands-on laboratory session. Participants completed lab activities using GPU-enabled VMs. They used Microsoft Azure Labs and Google Colab. Everyone had access to GitHub repositories containing lab notebooks, datasets, and code implementations. Participants worked with TensorFlow, PyTorch, differential privacy libraries, and federated learning frameworks. These tools represent standard implementations used in production AI systems.

Guest Speakers/Experts. The programs featured industry experts (*e.g.,* PayPal, Meta, and YSecurity) and academic researchers (*e.g.,* Loyola University Chicago, University of Wisconsin-Madison, North Carolina State University, Sungkyunkwan University, Arizona State University, University of California, Irvine, and Brown University). This breadth of expertise ensured participants received multiple perspectives on AI security and privacy challenges.
Program Runs. We conducted the first program run over six months (May-October 2024) with 12 bi-weekly workshops. We scheduled sessions from 6:00-8:00 PM to accommodate working professionals. We organized the curriculum following the 12-workshop structure.

The second program run occurred over six weeks (January-February 2025) with 6 weekly workshops. We maintained the 6:00-8:00 PM schedule. We condensed the curriculum by combining related topics into single workshops. This format combined pre-recorded lectures with live laboratory sessions to maximize learning efficiency. The program condensed the original 12 workshops into 6 sessions. Each covering multiple original workshop topics. For example, Workshop 1 combined *Introduction and Fundamentals* with *AI and Threat Models*. Workshop 2 merged *Adversarial Attacks* with *Robustness and Resilience*. This restructuring allowed for a shorter-length of time and a more intensive program while still addressing all key concepts.

C. Participant Recruitment and Characteristics

Recruitment occurred through multiple channels. These included industry advisory boards, the Center for Data Science Consulting, startup incubators, alumni networks, and social media. Target participants included AI practitioners, data scientists, software engineers, and security professionals working with AI systems. The first program recruited 26 participants, while the second attracted 30 participants. Both programs aimed for diverse professional backgrounds and experience levels to reflect real-world AI practitioner populations.

Figure 2 summarizes participant demographics and experience levels for both program runs. The first program had 26 participants enrolled in the full series; on the other hand, the second had 30 enrolled in the condensed format. The first column shows the summary for the first program, and the second column represents the summary for the second program. They are shown side by side to facilitate direct comparison between the two programs. Participants represented various sectors including industry, academia, government, and education. Experience levels ranged from less than one year to over ten years in AI-related roles.

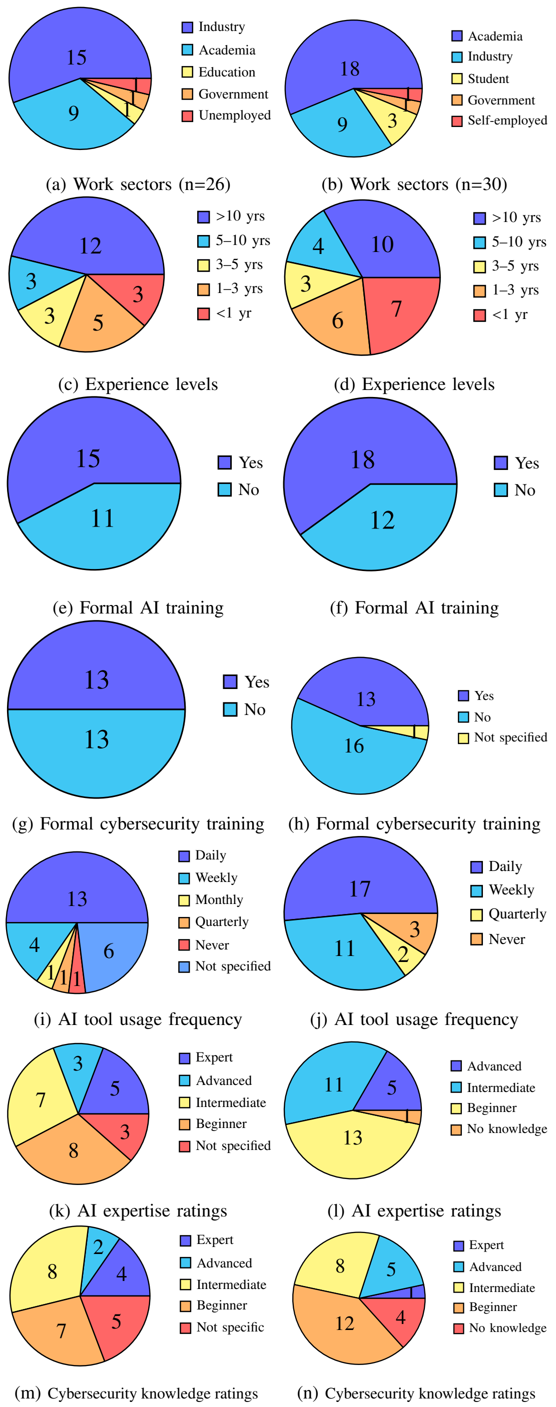

(a) Work sectors (n=26)
(b) Work sectors (n=30)
(c) Experience levels
(d) Experience levels
(e) Formal AI training
(f) Formal AI training
(g) Formal cybersecurity training
(h) Formal cybersecurity training
(i) AI tool usage frequency
(j) AI tool usage frequency
(k) AI expertise ratings
(l) AI expertise ratings
(m) Cybersecurity knowledge ratings
(n) Cybersecurity knowledge ratings

Fig. 2: Participant characteristics for first (left column) and second runs (right column).

Participant characteristics for both program runs as summarized as follows.

For the first run, participants (n=26) were primarily from industry (58%). They had significant representation from academia (35%) and smaller percentages from education, government, and unemployed sectors. Experience levels varied, with a notable portion having over 10 years of experience (46%). A majority had formal AI training (58%), while half had formal cybersecurity training (50%). AI tool usage was frequent, with 50% using them daily, 15% weekly, with remaining usage less frequent. Self-assessed expertise in AI ranged from expert (19%), 12% advanced, 27% intermediate, 31% beginner. Cybersecurity knowledge similarly distributed with 15% expert, 8% advanced, 31% intermediate, 27% beginner.

For the second run, participants (n=30) had higher academic representation (60%), followed by industry (30%), with students (10%), government (3%), and self-employed (3%). Experience levels were more diverse, with 33% having over 10 years of experience, and others distributed across experience ranges. A majority had formal AI training (60%), while 43% had formal cybersecurity training. AI tool usage was also frequent, with 57% using them daily, 37% weekly, and some using quarterly (7%) or never (10%). Self-assessed AI expertise showed no experts, with 17% advanced, 37% intermediate, 43% beginner, and 3% no knowledge. Cybersecurity knowledge showed 3% expert, 17% advanced, 27% intermediate, 40% beginner, and 13% no knowledge.

IV. Evaluation and Results

We conduct the evaluation with an independent evaluator for the program. The evaluator provides iterative analyses and actionable recommendations after each workshop series. This evaluation informs continuous curriculum improvement and assess program impact. The evaluation employs a mixed-methods evaluation approach including quantitative and qualitative assessments. Pre-program surveys assessed baseline knowledge in AI concepts, cybersecurity principles, and privacy-preserving techniques. Post-program surveys measured knowledge gains and skill development. Qualitative evaluation includes structured observations during laboratory sessions and analysis of participant projects and reflections. This approach provides transparent insights into learning processes and practical application challenges.

For the first run, five (n=5) participants completed the program, while for the second run, 11 participants (n=11) completed the program.

First Run Evaluation. Of the five respondents who completed the first workshop, two respondents completed a post-workshop survey. These two respondents differed on job experience. One assessed AI security threats several times a week. The next person never assessed AI threats as part of their job in the last year. Both respondents indicated that lectures were excellent. The clarity of directions for technical assignments was very good. The feedback on assignments was very good or excellent. Both respondents correctly answered 24 of the 30 knowledge questions on the post survey. One respondent commented "The guest presenter was great. I appreciated that they continued to tie back what was covered this week to past presentations, even though they weren't the one giving those presentations, both as a reminder of what we've discussed so far and to put what we're learning into context." Another appreciated the deeper understanding they gained of the robustness and designing robust models. The evaluator noted that across the workshop several students were highly engaged. They asked clarifying or critical questions. Meanwhile, they were able to contribute to group discussions about the exercises or lab assignments. The presenters provided clear, engaging talks that highlighted practical issues involving identifying and addressing cybersecurity threats in artificial intelligence tools.

Second Run Evaluation. The evaluation consisted of an online pretest (n=19) and an online post-test after the workshop was completed (n=5). It also included observations of each workshop. Regularly, there were 8 respondents with participation reduced after the first two workshops. Of those who completed the pre-test, one respondent attended only one of the sessions. For the five answering the post-test, two had been working on artificial intelligence projects for only two months, 1 for one to six months, 1 for one to two years, and one for six to seven years, indicating that the audience had a vast array of experience. Echoing this difference, respondents had a wide range of answers to how often have you assessed threats to the security of AI in the last year during your job? Answers included 'Never' (1), 'Very Rarely' (2), and 'Several times a week' (2).

Evaluator Notes. Institutional Review Board (IRB) approval was obtained before any survey data was collected. Each participant who completed the program and answered both the pre-test and post-test survey also received an Amazon gift card.

The second workshop on cybersecurity and privacy training program designed for AI professionals consisted of six workshops. They were offered once a week on Wednesday nights from 6:00 PM to 8:00 PM. Participants had access to readings and videos before each workshop. Meanwhile, they had access to the lab assignment that would be covered in the second half of the class. The content of the workshop was very organized. It provided critical information at a beginning level on cybersecurity and privacy issues in artificial intelligence. Some of the content is highlighted in this paragraph.

The first week provided an overview of critical concepts, providing a foundation for all future weeks and in a clear, engaging and informative presentation. Critical concepts included: Predictive vs. Generative AI models, the goals of adversarial attacks to target availability, integrity or confidentiality of data, the taxonomy of attacks on predictive and generative AI systems. It covered how attacks affect integrity, availability and privacy. Such as the types of adversarial attacks, threat models and examining robustness. It also covered possible defenses to implement against these attacks. The presentation covered timing of poisoning attacks to data and models during training and evasion and privacy attacks during deployment. It covered the capabilities of adversarial attacks and white-box, grey-box, and black-box attacks for predictive and generative AI tools.

The second module focused on privacy and how to balance privacy risks and utility. The third module focused on the critical concept of fairness and bias in AI systems and ways to mitigate bias and create models that prioritize fairness. The fourth module focused on transparency of AI systems. It discussed reasons for transparency such as regulatory accountability and enhancing user trustIt provided examples of poisoning and evasion attacks in models that could be used in self-driving cars, cybersecurity and medical imaging. It explained that transparency can assist in detecting and mitigating attacks. It also explained that transparency increases interpretability and explainability, providing easily understood examples.

The fifth week focused on what makes machine learning models vulnerable. It covered a variety of potential attacks such as cloning as well as the roles of cybersecurity professionals compared to machine learning researchers. The presentation examined how machine learning is used in cybersecurity operations. It examined the benefits and challenges of using ML, and how cybersecurity professionals view ML explanation techniques. The last week focused on real-world cases in AI security and privacy, summarizing the challenges, security and privacy principles, and best practices.

Evaluator Observations. The evaluator attended all 6 sessions and observed the lectures and lab assignment section. Each workshop consisted of a lecture and then a lab that was guided by an assistant. From the evaluator's perspective, the lectures were clear. They were all at a level that individuals who were not familiar with artificial intelligence could grasp the critical information. The information related to threats and attacks and ways to mitigate these threats and attacks in artificial intelligent tools. Workshop participants routinely asked questions. Small group discussions were used to enhance the understanding of critical issues. Several presentations also included questions to the audience to create engagement and check on understanding.

In the first two sessions, the participants were unprepared to use the lab tools. Some did not understand basic computer usage such as unzipping a file. This was unexpected as professionals who are not in computer science fields understand these basic commands. Future workshops need to decide whether to continue to offer the workshop to an audience with vastly different levels of computer and AI knowledge ranging from limited to advanced.

A. Results

Second Run Results. Table II presents a comparison of the pre-test and post-test samples on knowledge gained. As shown in Table II, the pre-test sample had a great deal of knowledge about basic cybersecurity or computer commands in Linux. Half of the sampled answered 87% of the 15 questions correctly. However, most participants had little knowledge about cybersecurity of artificial intelligence. Half answered 13% of the 15 questions correctly. In comparison, two of the five respondents who answered the post-test now answered 12 or 14 of the 15 questions correctly. The other three respondents answered 3, 6 or 7 of the same AI questions on the pre-test correctly.

Respondents on the post-test also were asked, "After attending the workshop, how much knowledge do you have about fundamentals of cybersecurity as it relates to Artificial Intelligence" on a 1 to 5. 1 indicated a little knowledge, 3 indicated moderate knowledge, and 5 indicated a great deal. Respondents reported moderate knowledge (Mean = 3.2, Median = 3.0). The responses ranged from 2 to 5.

Post-test Ratings from 5 respondents. Overall, respondents had a mixed review of the workshop. Table III provides three of the ratings. The respondents provided an overall 'good' average rating for the lectures and quality of feedback. None indicated needs improvement. Two of the five respondents indicated that it needs improvement for the clarity of lab assignments. The mean was 2.8. Respondents generally thought that they had the right amount of opportunity to discuss the class materials. Only one respondent indicating substantially too little. Respondents also thought that the group discussion moderately to greatly reinforced and added to the learning from the lectures. Similarly, respondents thought that technical lab assignments moderately to greatly reinforced the learning from lectures.

Only two respondents discussed strengths:

❶ *"The workshops: 1) provide an in-depth look about the security in AI 2) encompassed the techniques necessary for the security in AI 3) provide a theoretical as well as practical way for studying the security in AI."*

❷ *"Topics related to federated learning, case study, practical implementation."*

Only three respondents provided qualitative remarks about ways to Improve:

❶ *"Work on the tech of the labs. I enjoyed the privacy lab as it was a break from the dataset testing every week. The class members had their screens turned off and most likely were not there and it seemed disrespectful. If there is going to be a live class, screens on."*

❷ *"1) improve the labs for each workshop, 2) removing unnecessary code from files located in Github, 3) prepare a runtime environment to test the code, 4) improve the slide for each workshop, 5) in my case, I find a big problem to do the activities since colab is not free and to use more resources we should buy them such as GPU."*

❸ *"More interaction."*

B. Limitations and Future Work

This study's limitations include relatively small sample sizes and short-term follow-up periods. Longer-term assessment of knowledge retention and career impact would strengthen evidence for program effectiveness. The participant pool primarily consisted of technology professionals. This limits generalizability to other industries adopting AI systems. Healthcare, finance, and manufacturing sectors may require specialized adaptations. The evaluator observed a wide disparity in participant technical readiness. Some struggled with basic computer tasks such as unzipping files. Future programs should introduce an optional prerequisite assessment or a mandatory, non-credit foundational module. This would ensure all participants are equipped for the technical demands of hands-on labs. The program should also be extended to the future workforce. This includes undergraduate and graduate students who will enter AI-related fields.

Future work should explore personalized learning paths, and integration with existing professional development programs. Research into optimal combinations of synchronous and asynchronous learning components could improve program efficiency.

V. Conclusion

The experiential training program demonstrates that targeted, hands-on education can address the

TABLE II: Knowledge about Basic Cybersecurity and Artificial Intelligence Cybersecurity

Types of Knowledge	Pre-test Knowledge			Post-test Knowledge		
	Mean	Median	# of Items	Mean	Median	# of Items
About Basic Cybersecurity	73%	87%	15	85%	86%	7
About Artificial Intelligence	20%	13%	15	56%	47%	15

TABLE III: Post-test Ratings of Lectures and Lab assignments (n=5)

Concept	Mean	Median	Lowest score (n)	Highest score (n)
Quality of Lectures	3.2	3	2 (2)	5 (1)
Clarity of Lab Assignments	2.8	3	1 (2)	5 (1)
Quality of Feedback on Homework Assignments	3.0	3.0	1 (1)	5 (1)

Note. Scale is 1 to 5 with 1 = needs improvement, 2 = satisfactory, 3 = good, 4 = very good, 5 = excellent.

AI security knowledge gap among working professionals. Both comprehensive and condensed formats achieved significant learning outcomes, with the condensed approach showing better completion rates while maintaining educational effectiveness. The program format comparison reveals that intensive, focused training better accommodates professional schedules while maintaining learning outcomes. This finding has implications for designing scalable cybersecurity education programs that balance depth with accessibility. Our distinct contribution does not lie in being the first AI security training program. Rather, it lies in the structured, comparative evaluation. This evaluation compares long- versus short-format training delivery. We provide empirical evidence for program effectiveness. This evidence covers different formats. It offers guidance for future implementation decisions. Future work should explore long-term retention, career impact assessment, and adaptation to industry-specific contexts. The validated competency framework and assessment tools provide a foundation for broader implementation and continuous improvement.

Acknowledgment

This work was supported by the National Science Foundation under Award No. 2335700.

References

[1] W. Cui, T. Chen, C. Fields, J. Chen, A. Sierra, and E. Chan-Tin, "Revisiting assumptions for website fingerprinting attacks," in *ACM Asia Conference on Computer and Communications Security*, ser. AsiaCCS '19. ACM, 2019.

[2] G. Ho, A. Mirian, E. Luo, K. Tong, E. Lee, L. Liu, C. A. Longhurst, C. Dameff, S. Savage, and G. M. Voelker, " Understanding the Efficacy of Phishing Training in Practice ," in *2025 IEEE Symposium on Security and Privacy (SP)*. Los Alamitos, CA, USA: IEEE Computer Society, May 2025, pp. 37–54.

[3] L. Vaishnav, S. Singh, and K. A. Cornell, "Transparency, security, and workplace training & awareness in the age of generative ai," *arXiv preprint arXiv:2501.10389*, 2024.

[4] J. Prümmer, T. van Steen, and B. van den Berg, "Assessing the effect of cybersecurity training on end-users: a meta-analysis," *Computers & Security*, vol. 150, p. 104206, 2025.

[5] N. Al-Dhamari and N. Clarke, "Gpt-enabled cybersecurity training: A tailored approach for effective awareness," *arXiv preprint arXiv:2405.04138*, 2024.

[6] "Ai-powered adaptive cybersecurity awareness training for the industrial sector," *International Journal of Intelligent Systems and Applications in Engineering*, vol. 12, no. 4, pp. 5493–5505, June 2024, simulation-based study of 100 industrial employees demonstrating 72% reduction in phishing susceptibility.

[7] A. A. Nambiar, "Securing student data in the age of generative ai: A tool for data privacy enhancement in k12 schools," MIT Responsible AI for Social Empowerment and Education (RAISE), MPA Capstone Project, 2024.