Session 1: AI for Threat Detection & Synthetic Media Defense

November 12th at 10:40 AM

Deepfake-Enabled Infiltration: The Threat of Synthetic Identities in Corporate Environments

Joseph Lozada

This paper explores the evolving threat of deepfakes in the context of insider threats, particularly how advanced persistent threats (APTs) are leveraging AI-generated audio and video to impersonate job applicants and gain access to sensitive systems. While deepfakes have legitimate applications in entertainment, education, and business, they are increasingly being weaponized for deception and cyber intrusion. The paper outlines recent incidents, assesses technical vulnerabilities, and evaluates current risk management frameworks such as NIST RMF. It concludes with policy and technology recommendations to enhance detection and prevention strategies, especially during remote hiring and onboarding processes.

November 12th at 11:40 AM

Unequal Risks: Ethnicity, Region, and Cybersecurity Outcomes in the United States

Kaushik Reddy Mitta, Marc Dupuis

Cybersecurity risks are often treated as uniform, yet disparities across demographic groups suggest otherwise. This study investigates how ethnicity and geographic region shape cybersecurity outcomes in the United States, focusing on victimization, tool adoption, and awareness. Survey data from 470 adult participants were analyzed using ANOVA, Kruskal–Wallis tests, chi-square analyses, and logistic regression models. Results reveal two paradoxes. First, Asian/Pacific Islander respondents reported higher awareness and greater use of protective tools, yet also faced significantly elevated odds of identity theft, phishing losses, and account takeovers. Second, suburban residents exhibited higher preparedness than urban or rural populations, but consistently experienced greater exposure to cyber incidents, particularly financial fraud. Hispanic/Latinx and rural groups reported lower adoption of security tools, reflecting barriers of access and language. These findings highlight that awareness and adoption alone are insufficient when structural vulnerabilities and targeted exploitation are at play. The study underscores the need for culturally competent education, expanded infrastructure access, and adaptive monitoring systems to reduce disparities and promote a more equitable cybersecurity landscape.

Session 2: AI-Driven Operations & Convergent Security

November 12th at 2:00 PM

A Deweyan Foundation for Cultivating Reflective Cyber-Attuned Habits in an Age of AI and Ambiguity

Jane Blanken-Webb

Rapid advances in AI, automation, and hyperconnectivity are outpacing human habits, producing pervasive ambiguity. Drawing on John Dewey's philosophy of habit as growth through disruption and inquiry, this paper reconceptualizes cybersecurity education as cultivating reflective, cyber-attuned habits across society—not only training specialists. Dewey's account of growth through disruption, inquiry, and reorganization are translated into three educational design moves: (1) embed reflective inquiry within procedural exercises; (2) employ inquiry-based, experiential formats (e.g., capture-the-flag, cyber-defense exercises, cyber-ranges) to practice reasoning under uncertainty; and (3) extend learning to social practices of verification and shared deliberation beyond technical settings. The approach turns error into material for growth and equips learners to act with intelligent adaptability.

November 12th at 2:20 PM

AI-Driven Cloud Security: AIOps for Threat Detection and Compliance

Advait Patel, Vaishnavi Gudur, Prashanthi Matam, Charit Upadhyay, Aparna Achanta, Swara Dave, Shalini Sudarsan

The rapid growth of cloud and hybrid computing has brought significant scale, complexity, and security challenges to IT operations. Traditional rule-based monitoring systems and signature-based Security Information and Event Management (SIEM) tools are no longer sufficient to process the enormous volume of events generated in modern environments or to provide timely, accurate detection of incidents. Artificial Intelligence for IT Operations (AIOps) has emerged as a transformative approach by combining machine learning, predictive modeling, big data analytics, and automation to improve anomaly detection, optimize resource allocation, and accelerate the process of identifying root causes. Empirical studies report that AIOps platforms can reduce mean time to detection by nearly half and cut audit preparation time by up to 60%, underscoring their advantages over conventional methods. In addition to performance monitoring, AIOps is increasingly applied to security and compliance, enabling automated evidence collection, support for zero-trust architectures, and AI-assisted remediation workflows. Despite these benefits, reliance on opaque "black-box" models raises concerns around explainability, accountability, and regulatory compliance, particularly in mission-critical domains. Multi-cloud and hybrid infrastructures further complicate deployment due to interoperability issues, data silos, and risks of algorithmic bias. This paper reviews academic and industry work on AI-driven cloud security and operations from 2022 to 2025, outlines a taxonomy of AIOps functions spanning

detection, compliance, response, and governance, and identifies unresolved challenges such as adversarial resilience, transparency, and multi-cloud coordination. Finally, future directions are discussed, including explainable and neuro-symbolic AIOps, federated analytics for distributed environments, and autonomous self-healing infrastructures. The review aims to provide researchers and practitioners with a consolidated reference for developing trustworthy, scalable, and secure AI-driven cloud operations.

November 12th at 2:40 PM

Best Practices in Security Convergence: Tales from the Trenches

Michael Whitman, Herbert Mattord, Kathleen Kotwica

Security convergence, the integration of cybersecurity and physical security, has been discussed for over two decades, yet organizations still face challenges in defining and implementing effective strategies. This article explores the evolution of security convergence, highlighting key overlaps between physical and virtual security, especially with the rise of computerized operations in critical infrastructure. Through a 2023 survey and in-depth interviews with security professionals, four best practices are identified: implementing employee risk ratings, utilizing decision matrices, establishing fusion centers, and fostering a supportive organizational culture. These practices enhance collaboration and optimize security operations, demonstrating that effective convergence is not just about structural integration but also strategic coordination and cultural alignment for improved organizational resilience.

November 12th at 3:00 PM

EQiLevel: Emotion-Aware Reinforcement Learning for Multimodal Academic Tutoring

Veronica Elze, Taejin Kim, Jooyol Maeng

EQiLevel is an emotionally adaptive AI tutoring system that integrates reinforcement learning (RL), large language models (LLMs), and real-time sentiment detection to personalize instruction dynamically. Traditional intelligent tutoring systems often follow rigid rules and fail to account for learners' emotional states or individual learning preferences, reducing engagement. EQiLevel addresses this limitation by analyzing voice-based cues and adapting lesson difficulty, tone, and pacing in real time through a JSON-based Model Context Protocol (MCP). The MCP encodes emotion, performance, and learning style variables into structured state information, guiding GPT dialogue generation and reinforcement learning policy updates. Evaluation with simulated data demonstrated 78% successful adaptation to frustration cases, Whisper transcription accuracy with a 5.3% word error rate, emotion detection accuracy of 84% with 81% tone alignment, and improved RL convergence with average rewards rising from 0.41 to 0.63. In the context of cybersecurity education, EQiLevel illustrates how adaptive, emotion-aware tutoring can prepare learners to remain resilient under ambiguous and adversarial conditions, such as phishing awareness and threat analysis. By uniting technical adaptability with emotional intelligence, EQiLevel provides a scalable framework for inclusive, resilient, and effective cybersecurity education in the age of AI and automation.

November 12th at 3:20 PM

Self-Hosted Workflow Automation For AI-Based Cybersecurity Operations

Hareign Casaclang, Bianca Ionescu, Yoohwan Kim, Ju-Yeon Jo

Cybersecurity operations often involve repetitive tasks such as running Nmap scans, analyzing logs, and performing Open-Source Intelligence (OSINT) investigations. These processes are essential for maintaining security but consume time and resources that many small organizations cannot spare. While commercial automation platforms exist to reduce this workload, they are typically costly and inaccessible to businesses without dedicated IT staff. This paper investigates n8n, a self-hosted and low-cost workflow automation platform, as a practical alternative for cybersecurity automation. By integrating security tools and external large language models (LLMs) such as ChatGPT, Gemini, and Ollama, n8n can automate vulnerability scanning, assign severity ratings, and generate reports tailored to both technical and executive stakeholders. Experiments show that n8n workflows can effectively combine traditional scans with Artificial Intelligence (AI)-driven analysis to produce actionable outputs. Although limitations remain, including a steep learning curve and restrictions in the free tier, n8n demonstrates potential for broadening access to automation in cybersecurity. For small organizations, this approach provides a cost-effective way to strengthen security posture, while in academic contexts it provides a hands-on platform for teaching and experimenting with automation and AI in cybersecurity.

Session 3: Cyber Risk Awareness & Human Behavior

November 12th at 4:00 PM

Analysis of Cybersecurity Risks and Teenage Digital Behavior Patterns

Eric McCloy, Samuel Nimako-Mensah, Albert Samigullin

As teenagers increasingly engage with digital technology, cybersecurity vulnerabilities present significant risks to their online safety and privacy. Adolescents who lack awareness of secure online practices are particularly vulnerable to malicious actors seeking to exploit them. This paper investigates the relationship between real-world online behavior of teenagers, cybersecurity risks, and device interactions. The primary data set used for this analysis is Teenage online behavior and cybersecurity risks. First, we consider demographic information: age, education, time spent online. We correlate this with online behaviors: use of a VPN, type of equipment (computer, mobile), use of public internet, engagement with risky websites. Finally, we analyze the data set using a combination of demographic and behavioral patterns to search for high-risk, negative outcomes. For our research, we analyze teenage online behavior patterns to identify key risk factors, develop predictive models for cybersecurity vulnerabilities, and produce actionable visualizations that illustrate the relationship between digital literacy and online safety. Our findings utilize data and business analytics to provide evidence-based recommendations for parents, educators, and policymakers to enhance teenage cybersecurity awareness and protective strategies.

November 12th at 4:20 PM

CodeWars: Using LLMs for Vulnerability Analysis in Cybersecurity Education

Arunima Chaudhary, Walter Colombo, Amir Javed, Junaid Haseeb, Vimal Kumar, Fernando Alva Manchego, Richard Larsen

Large Language Models (LLMs) are increasingly explored as tools for software development and could further constitute a supplementary source for the development of varied examples intended for pedagogical use. While they can improve productivity, their ability to produce code that is both secure and compliant with Secure Software Development (SSD) practices remains uncertain, raising concerns about their role in cybersecurity education. If LLMs are to be integrated effectively, students must be trained to critically evaluate generated code for correctness and vulnerabilities, raising an important question: How can LLM-generated code be effectively and securely incorporated into Cybersecurity education for teaching vulnerability analysis? This paper introduces CodeWars, a novel teaching methodology that combines LLM-generated and human-written code to examine how students engage with vulnerability detection tasks. CodeWars was implemented as a pilot study with a total of 32 students at Cardiff University and the University of Waikato, where students analyzed flawed, secure, and mixed-origin code samples. By comparing student approaches, analysis, and perceptions, the study provides insights into how vulnerabilities are detected, how code origins are distinguished, and how SSD practices are applied. Our analysis of student feedback and interviews indicates that Codewars produced structured and accessible code, simplifying vulnerability identification and offering educators the means to efficiently develop varied SSD teaching

applications. These findings illuminate both the advantages and constraints of employing LLMs in secure coding and position this study as a foundational step toward the responsible adoption of AI in Cybersecurity education.

November 12th at 4:40 PM

From Social Sharing to Security Lessons: Behaviors, Disclosure, and Cyber Threats

Sav Wheeler, Marc Dupuis

As social networking sites (SNSs) have become integral to daily life, concerns about privacy and cybersecurity risks have intensified. Malicious actors exploit SNSs for phishing, malware distribution, and identity-driven attacks, often leveraging personal information voluntarily disclosed by users. This study investigates the relationships between SNS usage, personal information disclosure, cybersecurity behaviors, and experiences with cybersecurity threats. We employed a mixed-methods approach, combining survey data from 275 participants with semi-structured interviews. Correlation analyses revealed that frequency of SNS use and usage motivations---particularly for meeting new people and for self-presentation---were positively associated with higher levels of personal information disclosure. Disclosure of personal information and frequency of SNS usage were also significantly correlated with reported experiences of cybersecurity threats, though less so with protective cybersecurity behaviors. Interview responses highlighted both direct encounters with threats and broader perceptions of privacy vulnerabilities. Together, these findings underscore the complex interplay between social behavior on SNSs and cybersecurity risks, suggesting that greater user education and platform-level safeguards are necessary to mitigate emerging threats. We conclude with implications for cybersecurity awareness efforts and recommendations for future research.

November 12th at 5:00 PM

Past Experience and Threat Awareness as Determinants of Regular Information Backup

Benard Birundu, Marc Dupuis

Cybersecurity threats continue to evolve in sophistication, increasingly targeting individuals as the weakest link in the security chain. While technical solutions remain essential, human-centered protective behaviors such as regular information backup are critical to mitigating risks from ransomware, device failure, and accidental deletion. This study investigates how past experiences with cyber incidents and awareness of threats influence backup practices among individuals in the United States. Drawing on Protection Motivation Theory (PMT), we conducted a mixed-methods survey (N{=}308) that measured threat appraisal (perceived severity, vulnerability) and coping appraisal (self-efficacy, response efficacy, response cost), along with threat awareness and prior experience. Multiple regression explained nearly half of the variance in backup behavior ($R^2{=}0.498$), with self-efficacy and threat awareness emerging as strong positive predictors; response cost was negative and significant. Qualitative responses illustrated experience- and awareness-driven adoption and highlighted common hybrid routines (cloud plus removable media) as well as barriers related to perceived effort or low perceived value. The findings

underscore the importance of integrating human factors into cybersecurity programs and suggest concrete levers for awareness, design, and policy that reduce friction and promote routine backups.

November 12th at 5:20 PM

Roll with it: Awareness raising with Cyber Defence Dice

Steven Furnell, Lucija Šmid, James Todd, Xavier Carpent, Simon Castle-Green

Cybersecurity awareness is widely recognised as an important requirement, but is frequently overlooked or addressed in ways that do not engage the interest of the target audience. In an attempt to broaden the options available for achieving this, the paper discusses the concept, design and evaluation of a new dice-based game designed to promote entry-level cyber security awareness in relation to common forms of attack and defence. The concept of the game is that players can defend against prior attacks, or use attacks to test defences, with the images on the different die faces denoting threats and safeguards of different strengths and which are countered in different ways. The discussion includes a worked example of one of the game modes that has been designed in order to illustrate how players would take turns and make decisions in practice. It then presents initial results from a series of seven hands-on playtest sessions conducted with a range of audiences, including the general public, cyber educators and cyber professionals. The findings indicate that all audiences were positive about the game concept and found it enjoyable to play. Additionally, it was recognised to have value in raising and maintaining awareness, and would be a game that participants would play again and recommend to others.

Session 4: Cybersecurity Workforce Development & Frameworks

November 13th at 10:20 AM

Mapping the Gap: Analysis of Nuclear Cybersecurity Education in U.S. Universities

Myles Nelson, Amorita Christian, Tiffany Fuhrmann, Charles Nickerson

The U.S. nuclear sector is undergoing rapid transformation, driven by the expansion of advanced reactors, digital modernization of legacy systems, and increasing interest in nuclear energy to meet AI-fueled energy demands. However, the cybersecurity talent pipeline is not keeping pace with this growth. This paper investigates the significant gap in nuclear cybersecurity education and proposes scalable strategies for colleges to address this critical need by promoting it as a viable and essential career path.

Through a multi-institutional landscape analysis of 16 cybersecurity and 12 nuclear engineering programs, we found that nuclear cybersecurity is largely absent from university curricula. Most students are unaware of the field's existence, and few institutions offer hands-on training or interdisciplinary exposure. This lack of awareness leads to a shortage of specialized talent, forcing nuclear facilities to retrain generalist hires or rely on costly external consultants.

We present a framework for early pipeline cultivation grounded in Social Cognitive Career Theory and workforce development principles. Proposed solutions include student-led clubs, guest lectures, modular classroom kits, and summer boot camps. By increasing visibility and access to nuclear cyber content, we aim to break the self-reinforcing cycle of low awareness and limited specialization. This work underscores the critical role of education and advocacy in cultivating early interest and guiding students toward this emerging field. We call on academic institutions, national laboratories, and industry stakeholders to collaborate in establishing nuclear cybersecurity as a distinct and accessible career path within the broader cybersecurity and nuclear engineering ecosystems.

November 12th at 10:40 AM

Lost in Translation: Evaluating Cybersecurity Policy and Terminology Accuracy for Multilingual Learners

Andrew Hurd, Gloria Kramer, Pamela Doran

As cybersecurity becomes a cornerstone of global higher education, language has emerged as an unexpected point of vulnerability. Machine translation (MT) tools, increasingly used to render cybersecurity policies into multiple languages, often distort meaning by translating technical terms literally rather than conceptually. Words like firewall, phishing, or backdoor lose their intent in translation, creating barriers to comprehension and leaving multilingual learners at risk of misunderstanding critical policies. This paper explores the intersection of cybersecurity vocabulary, machine translation, and language equity, drawing on examples of mistranslations and language acquisition research to demonstrate how linguistic gaps can weaken both institutional safeguards and student confidence.

We argue that cybersecurity education must treat terminology with the same precision as code, recognizing that mistranslation not only undermines

clarity but also compounds anxiety for multilingual learners navigating complex technical content. To address these challenges, the paper examines strategies such as the use of back-translation, custom glossaries, Universal Design for Learning (UDL) frameworks, and emerging AI translation tools like custom ChatGPT models. Together, these approaches highlight a pathway for higher education to balance inclusion with accuracy, ensuring that policies and coursework maintain both technical rigor and accessibility. By reframing cybersecurity not only as a technical field but also as a linguistic one, this research calls for a more intentional, equity-driven approach to translation that secures both data and learning outcomes.

November 13th at 11:00 AM

Building Nuclear-Specific Cybersecurity Expertise in Higher Education

Amorita Christian, Myles Nelson, Tiffany Fuhrmann

The rapid digitalization of nuclear power plants (NPPs) and the deployment of advanced and small modular reactors (A/SMRs) have expanded the cybersecurity attack surface within the nuclear sector. This evolution introduces unique challenges beyond those faced in general information technology (IT), operational technology (OT) and industrial control system (ICS) security, due to nuclear power's regulatory rigor, safety-critical nature, and operational needs.

A pressing workforce gap persists; cybersecurity graduates typically lack nuclear-specific context and retraining them for industry readiness requires 12–18 months, creating a significant burden. This paper addresses this gap by defining the domains of knowledge that nuclear cybersecurity specialists must master, spanning cybersecurity, nuclear engineering, OT/ICS security, and regulatory governance. We propose a curricular framework integrating technical, regulatory, and applied learning components to accelerate workforce readiness.

Our approach builds on existing findings that current curricula inadequately integrate nuclear engineering and cybersecurity, shifting the discourse from why specialization is needed to what knowledge must be taught. The recommendations have implications for workforce development and long-term resilience of the nuclear energy sector.

November 13th at 11:20 AM

Play NICE: Incorporating Cyber Phraseology into K-12 Education

Timothy Crisp, John Hale

Cyberattacks on critical infrastructures motivate a focus on cybersecurity awareness. A knowledge gap exists in the technical and non-technical understanding of cybersecurity, in the workforce. Closing this gap requires a multi-faceted approach -- of extreme importance is education. We use the NICE Workforce Framework TKS statements to develop a model of the most generalizable requirements needed to practice cybersecurity. We apply this model to increase cybersecurity language use and comprehension in all K-12 subjects, walking educators through a process incorporating cybersecurity into their lessons.

November 13th at 11:20 AM

Teaching Critical Infrastructure Security Through Interactive Experiences: Modeling Cyberattacks in Gamified Learning

Ella Luedeke, Meera Sridhar, Harini Ramaprasad

This work introduces InfraLearn, a gamified learning platform designed to teach non-computer science students foundational background in cybersecurity for critical infrastructure. InfraLearn simulates attacks on a Distributed Energy Resource (DER) device, modeled after the Enphase Gateway solar monitor and implemented using a Flask-based API. Three prototype scenarios are developed: API spoofing, unauthorized remote shutdowns, and Living-off-the-Land (LoTL) downgrade exploitation. These scenarios are derived from real-world vulnerabilities in DER systems and integrated into a narrative-driven, web-based platform. Students interact with pre-configured virtual machines, guided code templates, and checkpoint quizzes, with optional AI support that reinforces comprehension while minimizing the need for prior programming experience. By situating cybersecurity concepts within the context of energy systems, InfraLearn makes abstract threats tangible and emphasizes the ethical application of defensive skills. This work demonstrates a scalable approach to engaging future engineers in securing critical infrastructure.

Session 5: Phishing, Awareness & Applied Security Training

November 13th at 2:00 PM

A Systematic Review of Residual Risk in Cybersecurity Awareness Training

Venkat Laxmi Sateesh Nutulapati

Cybersecurity awareness training is central to education and practice, yet persistent human error continues to expose organizations to breaches. AI-enabled attacks such as deepfakes, voice-cloned vishing, and automated spear phishing make these vulnerabilities even more consequential. This systematic review synthesizes 26 studies (2008–2025) using varied designs and training formats, from gamified learning and face-to-face sessions to e-learning, nudges, and simulated phishing. We introduced a residual-risk framework to capture outcomes that traditional effectiveness measures overlook. Residual Insecure Behavior (RIB) reflects the percentage of participants who continued risky practices after training, while Residual Knowledge Gap (RKG) indicates knowledge deficits that persisted. Across studies, improvements were common, but residual risks remained significant with phishing susceptibility often exceeding 10%, and knowledge gaps frequently surpassing 30%. Gamified approaches showed stronger behavioral effects, while conventional methods often raised awareness but left large gaps. For educators, these findings underscore that statistical gains can mask enduring weaknesses. By teaching and applying RIB and RKG, instructors can help students, practitioners, and organizations focus not just on learning outcomes, but on reducing real-world exposure in an AI-driven threat landscape.

November 12th at 2:40 PM

Detecting and Mitigating AI Prompt Injection Attacks in Large Language Models (LLMs)

Abel Ureste, Hyungbae Park, Tamirat Abegaz

AI is being interconnected with vital systems at an exponential rate, being described as the greatest shift in technology since the invention of the Internet. However, with the emergence of AI also involves the introduction of new critical vulnerabilities in the technology sector. This research will discuss the types of prompt injection attacks that AI can be subjected to, what they target and the possible repercussions of prompt injections. To counteract these attacks, solutions to detect different types of prompt injection will also be discussed, giving solutions to mitigate attacks that can expose critical data. Along with the solutions, different trade-offs between these solutions will be given. This research aims to expose the security issues that arise with the rapid implementation of experimental AI involving prompt injection and how to prevent it.

November 13th at 3:00 PM

Development and Validation of a Healthcare Workers Phishing Risk Exposure (HWPRE) Taxonomy for Mobile Email

Christopher Collins, Yair Levy, Gregory Simco, Ling Wang

Email on mobile has become a dominant communication channel for healthcare professionals, yet its constrained interface and context of use amplify vulnerability to social engineering attacks, especially phishing. This paper reports the

development and empirical validation of the Healthcare Workers Phishing Risk Exposure (HWPRE) taxonomy, a 2×2 framework that positions individuals by (i) general email-phishing susceptibility; and (ii) ability to detect mobile-specific phishing cues. We followed a sequential three-phase design: (1) a Delphi study with cybersecurity subject matter experts to validate mobile-relevant phishing indicators and components of a susceptibility index; (2) a pilot to refine instruments and procedures; as well as (3) a large-scale study (N=300 healthcare workers) using scenario-based assessments on smartphone-generated email stimuli. We present the construction of the Healthcare Workers Email Phishing Susceptibility Index (HWEPSI), reliability/validity evidence, and statistical analyses relating HWPRE placement to role, experience, medical departments, prior training, and demographic indicators. The results show significant heterogeneity across departments and experience bands; in addition, the ability to recognize mobile cues does not follow uniformly with general susceptibility. We discuss implications for targeted Security Education, Training, and Awareness (SETA) programs and measurement-driven program evaluation. We conclude with practical guidance for integrating HWPRE into organizational phishing defense and directions for future research.

Session 6: CCERP

November 13th at 4:00 PM

AI and ML Attacks on IC Hardware Security: Demonstration for Cybersecurity Students

Danai Chasaki

The IC design and security industry depends on trusted systems, yet remains challenged by an increasingly fragmented supply chain and evolving threat landscape. The rise of fabless enterprises and the proliferation of AI/ML technologies have further exposed hardware security to new vulnerabilities. This paper provides proof-of-concept implementations of emerging threats posed by machine learning to IC hardware design, focusing on two distinct areas: GNN-based attacks on logic locking and the insertion of hardware Trojans via large language models. These represent growing and independent research directions in hardware security. We showcase and analyze two representative examples from each category to highlight the risks of unmitigated ML-driven attacks.

November 13th at 4:20 PM

Cybercamp: An Experience Report on the Transformations of an Intensive Cybersecurity Summer Camp for High School Students

Jose R. Ortiz Ubarri, Rio Piedras, Rafael A. Arce Nazario, Rio Piedras, Kariluz Dávila Diaz

The Cybercamp is a Cybersecurity summer camp for high school students that has been held for the last nine years at a Hispanic Serving Institution. Since its inception in 2016 the Cybercamp has undergone several transformations in response to budget reductions and the COVID pandemic, to finally become its current version: a rich, hands-on learning experience that we believe is easily replicable even in resource challenged environments.

In this paper, we document the transformations of the Cybercamp and discuss the developed curriculum and materials in hopes that others will reuse, adapt, and improve upon them. In the Cybercamp, we apply active learning practices that have been shown to be effective in STEM education. We perform handson, activities, providing Capture The Flag (CTF) style practice exercises with automatic grading. The Cybercamp is assisted by college students who serve as peer-assisted leaders. The educational materials were designed with culturally relevant case studies, using open source technologies, and following Universal Design Learning best practices to make them as accessible as possible, particularly to low-income students. We discuss how we prevented students from falling behind and successfully completed the Cybercamp.

November 13th at 4:40 PM

Examining the Capabilities of GhidraMCP and Claude LLM for Reverse Engineering Malware

Joshua Cole Watson, Bryson R. Payne, Denise McWilliams, Mingyuan Yan

Can Large Language Models (LLMs) enhance the efficiency of reverse engineering by assisting malware analysts in the de-obfuscation of ransomware and other forms of malicious software? This research explores the integration of LLMs into reverse

engineering workflows through the use of GhidraMCP, a plugin designed for the Ghidra open-source software reverse engineering suite. GhidraMCP leverages the capabilities of Claude's Sonnet 4 model (as well as other LLMs) to rename decompiled variables and functions, generate descriptive annotations for disassembled code, and highlight potentially relevant strings or routines. These features are intended to reduce the cognitive load on analysts and accelerate the identification of critical components such as encryption routines, embedded URLs, command-and-control (C2) indicators, and external library calls within malware samples.

This study compares traditional reverse engineering workflows with LLM-augmented workflows using GhidraMCP. Multiple pseudo-ransomware samples were analyzed to assess differences in discovery efficiency, accuracy of function labeling, and qualitative analytical quality. Although no formal timing metrics were recorded, the research team determined that the LLM-augmented process consistently achieved insights more quickly and with fewer manual steps. In several instances, Claude Sonnet 4 successfully identified static relationships and artifacts that human analysts initially overlooked, demonstrating its potential to enhance traditional workflows through contextual inference and advanced pattern recognition.

The combination of GhidraMCP and Claude Sonnet 4 effectively leveraged static analysis to identify the hidden flags for ESCALATE challenges one through seven. However, while the research team was ultimately able to solve all challenges, several required dynamic analysis and binary patching—tasks that the current LLM-augmented setup could not perform due to the lack of patching capabilities within GhidraMCP. It remains unclear whether this limitation stems from Ghidra, the plugin, or the integration framework itself. During testing, Claude Sonnet 4 occasionally exhibited hallucinations, producing inaccurate or speculative annotations that required human correction and additional prompting, particularly during challenges three and four. These occurrences emphasize the ongoing need for human oversight and iterative validation when employing generative AI in critical cybersecurity tasks.

Despite these limitations, the findings indicate that LLM-augmented reverse engineering can meaningfully improve analytical comprehension, efficiency, and context awareness. Claude Sonnet 4's linguistic reasoning and ability to infer code intent proved especially valuable for de-obfuscating complex binaries. Future work will focus on enabling dynamic capabilities within GhidraMCP to support patching and execution-based testing, as well as refining prompt strategies and hallucination detection. This research establishes a foundation for the continued development of intelligent, LLM-assisted tooling designed to augment human expertise in malware analysis and reverse engineering.

November 13th at 5:00 PM

Integrating Vulnerability Assessments with Security Control Compliance

Ernesto Ortiz, Aurora Duskin, Noah Hassett, Clemente Izurieta, Ann Marie Reinhold

Information technology providers must implement security controls to protect client and partner data, as well as comply with government security requirements. The preparation of security compliance documentation is a slow process due to the manual

efforts involved. We present SSP Manager, a framework that streamlines compliance and supports the building and maintenance of System Security Plans. The tool integrates vulnerability assessment, program analysis, and control monitoring by i) implementing a security control prioritization strategy that outputs NIST SP 800-53 controls to mitigate MITRE ATT&CK techniques, ii) by incorporating reachability analysis of Python dependencies to filter out false positives from vulnerability scan results, and iii) by providing compliance monitoring of functionality based on Chef's InSpec testing framework and the Open Policy Agent policy engine. Test reports are generated in a machine-readable format for easy integration into automated compliance pipelines. Our work bridges the gap between vulnerability assessment and security compliance. Moreover, it reduces the manual overhead in security workflows.

Session 7: Generative AI in Cybersecurity Education

November 13th at 4:00 PM

Cybersecurity Education with Generative AI: Creating Interactive Labs from Microelectronic Fundamentals to IoT Security Exploitation

Kushal Badal, Xiaohong Yuan, Huirong Fu, Darrin Hanna, Jason Gorski

Creating engaging cybersecurity education materials typically requires months of development time and specialized expertise. This paper describes how we used generative AI to address this challenge. We utilized Claude AI to generate a complete interactive platform that teaches students basic microelectronics through IoT hacking. Through iterative prompting, we generated more than 15,000 lines of functional code, including interactive visualizations, Python security tools, and gamified quizzes with real-time leaderboards. The curriculum guides students through the evolution of computing—from vacuum tubes to modern IoT devices—then helps them apply this foundation to discover real vulnerabilities. We implemented this platform at a GenCyber summer camp with 40 participants, where students identified actual security issues in AmpliPi audio systems—open-source network audio devices designed for multi-room audio distribution—including password weaknesses and denial of service flaws. The entire development process took only three weeks instead of the typical several months. The AI produced quality educational content, although we reviewed everything for technical accuracy and ethical considerations. During the camp, students remained engaged through competitive elements and hands-on labs, learning both theoretical concepts and practical skills. The students used AI-generated tools, including working implementations of SlowLoris and dictionary attacks, to test real systems. Our experience demonstrates that generative AI can efficiently create effective cybersecurity education materials that remain technically current. All materials are publicly available on GitHub for educational use. This approach could help educators stay on track with the rapidly evolving technology despite traditional curriculum development constraints.

November 13th at 4:20 PM

Distributed Agency in AI-Enhanced Cybersecurity Education: A Posthuman Instructional Design Framework

Ryan Straight, Josh Herron

This paper addresses a critical challenge facing cybersecurity educators: preparing students for AI-enhanced practice environments where effective action emerges from human-AI collaboration rather than individual expertise. Traditional instructional design frameworks assume human-centered learning processes that inadequately address distributed agency realities in contemporary cybersecurity operations. Drawing on Adams and Thompson's posthuman inquiry methodology, this analysis develops a comprehensive pedagogical framework consisting of four principles: (1) Design for the Assemblage, Not the Individual, (2) Cultivate Relationality and Response-ability, (3) Embrace Emergence, Messiness, and Indeterminacy, and (4) Posthuman Assessment Approaches. The framework provides concrete instructional design implications, including strategies for configuring human-AI learning

relations, integrating AI literacies across cognitive, civic, creative, and critical dimensions, and developing assemblage-aware cybersecurity case studies. These design implications bridge theoretical posthuman concepts with practical curriculum implementation through the lens of curriculum-as-lived rather than curriculum-as-plan. Preliminary implementation observations from an undergraduate cybersecurity ethics course demonstrate how posthuman-designed scenarios enable students to develop comfort with complexity and distributed analysis. Student reflections reveal progression from seeking singular solutions to embracing multiple valid perspectives, suggesting effective cultivation of human-AI collaborative competencies. The framework equips cybersecurity educators with both theoretical foundations and actionable design strategies for preparing students for distributed agency practice environments.

November 13th at 4:40 PM

Integrating Artificial Intelligence into Undergraduate Cybersecurity Education: A Course Design for Threat Detection, Explainability, and Ethical Resilience

Vahid Heydari, Kofi Nyarko

This paper introduces an undergraduate course, Artificial Intelligence Applications in Cybersecurity, designed to equip students with Artificial Intelligence (AI) and Machine Learning (ML) skills to address modern cyber threats. The curriculum integrates supervised and unsupervised learning, deep learning, explainable AI (XAI), adversarial ML, and ethical considerations. Using accessible tools (Python, Google Colab) and real-world datasets (e.g., NSL-KDD, CICIDS2017, malware corpora), students complete phased projects progressing from classical ML baselines to deep learning with interpretability (SHAP/LIME) and robustness against adversarial attacks (FGSM/PGD with mitigation). The course aligns with data science and cybersecurity workforce frameworks, emphasizing reproducibility, communication, and responsible AI practices.

November 13th at 5:00 PM

Toward Experiential Training Program for AI Security and Privacy Practitioners

Mohammed Abuhamad, Mujtaba Nazari, Loretta Stalans, Eric Chan-Tin

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.

November 12th at 11:40 AM

An AI Agent Workflow for Generating Contextual Cybersecurity Hints

Hsiao-An Wang, Joshua Goldberg, Audrey Fruean, Zixuan Zou, Ruoyu Zhao, Sakib Miazi, Jens Mache, Ishan Abraham, Taylor Wolff, Jack Cook, Richard Weiss

Large Language Models (LLMs) have proven beneficial in aiding student learning across a multitude of domains such as computer science, data science, and mathematics. While these chatbots show promise, they can be impractical to deploy in situations where quality student data is not available.

Hint generation for cybersecurity can be feasible with the technology we have today if we make the right simplifications. First, we need human-in-the-loop systems because the training data used to train general LLMs may not cover cybersecurity well. Second, using modular agents allows us to access dynamic data that is outside the training dataset and add specificity to the hints.

In this paper, we leverage n8n, an agent deployment service, to establish the connections between our agents and Discord, the messaging system that our classroom uses to offer students a streamlined learning experience when working on interactive cybersecurity exercises. We have tested this in the classroom and have shown that it is feasible.

Session 8: CCERP

November 13th at 4:00 PM

Security and Privacy of Wearable and Implantable Medical Devices: A Health Informatics Course on Security and Privacy of Wearable and Implantable Medical Devices

Michelle Mical Ramim

As wearable and implantable medical devices become fundamental to remote patient monitoring and precision medicine, the associated security and privacy risks demand urgent attention. These devices are increasingly targeted by cybercriminals, potentially endangering patient safety and data integrity. Specifically, it has been documented that vulnerabilities in medical devices have been exploited to alter device behavior or interfere with clinical treatment delivery. Despite these known vulnerabilities, wearable and implantable medical devices have become integral to modern patient care, offering innovative ways to monitor, manage, and even remotely treat various health conditions. These devices are essential to remote patient monitoring and precision medicine; the real-time data they capture is increasingly integrated into electronic health records (EHRs) to support clinical decision-making and enhance workflow efficiency. At the same time, most medical and healthcare students around the nation are not well educated to deal with cybersecurity issues. Subsequently, medical and healthcare students should understand the vulnerabilities associated with wearable and implantable devices, the risks they pose, and the importance of regulatory compliance, including the Health Insurance Portability and Accountability Act (HIPAA). To address this, we developed an experiential learning course titled Security and Privacy of Wearable and Implantable Medical Devices, designed for advanced undergraduate and graduate students in health and medical fields. The course immerses students in real-world challenges through lectures, labs, and project-based learning, leveraging wearable devices such as FitBitTM to analyze and interpret real-time personal health data. The curriculum covers critical topics including data security, privacy, HIPAA compliance, data visualization, interoperability, and real-world cyberattack case studies. The learning objectives align with the Commission on Accreditation for Health Informatics & Information Management Education (CAHIIM) standards and Miller's Pyramid of Clinical Competence to ensure industry-relevant competencies and progressive skill development. Interactive lectures were designed to promote engagement and featured expert guest speakers from health information technology (IT) and cybersecurity sectors. Case-based discussions encouraged students to consider the implications of cyberattacks on patient safety and health outcomes. The lab component offered a structured environment for technical practice, such as configuring wearable devices, extracting and visualizing data, and evaluating the security of data transmission. Lab assignments played a central role in reinforcing the key concepts introduced in lectures and assigned readings. By combining didactic instruction with applied learning and real-world examples, the class components provided a robust experiential learning that mirrors current challenges faced by healthcare.

November 13th at 4:20 PM

Validating the fundamental cybersecurity competency index (FCCI) through expert evaluation for human-generative artificial intelligence (GenAI) teaming

Witko, Yair Levy, Catherine Neubauer, Gregory Simco, Laurie Dringus, Melissa Carlton

The increasing volume of cyber threats, combined with a critical shortage of skilled professionals and rising burnout among practitioners, highlights the urgent need for innovative solutions in cybersecurity operations. Generative Artificial Intelligence (GenAI) offers promising potential to augment human analysts in cybersecurity, but its integration requires rigorous validation of the fundamental competencies that enable effective collaboration of human-GenAI teams. Fundamental cybersecurity competencies, encompassing essential cybersecurity Knowledge, Skills, and Tasks completion (KSTs). Competency is defined as the ability to complete tasks within a work role. In this research study, we employed a mixed-methods research approach designed to evaluate human-GenAI teams, emphasizing the role of expert consensus in shaping the experimental assessment of the Fundamental Cybersecurity Competency Index (FCCI) in a commercial cyber range. Selecting a commercial cyber range allowed us to identify the specific KSTs from the United States (U.S.) Department of Defense (DoD) Cyber Workforce Framework (DCWF) and measure them at the KSTs level. The specific commercial cyber range we assessed enables the extraction of the users' performance at the KST level. To validate the proposed experimental assessment of the FCCI and confirm the relevance of the selected cybersecurity KSTs, a panel of 20 Subject Matter Experts (SMEs) was engaged to evaluate and validate the proposed competency measures. The expert panel refined the cybersecurity scenarios and experimental procedures used in the commercial cyber range hands-on scenarios, ensuring alignment with DCWF. Our findings indicated that 46 of 47 fundamental cybersecurity KSTs were validated by the SMEs as essential components of the FCCI. Consensus levels of 85–90% confirmed strong expert support for incorporating GenAI (e.g., large language models such as ChatGPT) as a teammate or decision-support agent in these controlled experiments. The validated scenarios and experiments pave the way for future research on assessing cybersecurity competencies in commercial cyber range platforms with and without GenAI support (e.g., large language models such as ChatGPT). By establishing the baseline for competency assessment in this research, the SMEs' feedback contributed to advancing cybersecurity workforce development and provided critical insights for integrating GenAI into collaborative cybersecurity human-GenAI teaming operations. The validated FCCI provides a robust mechanism to evaluate both human and human–GenAI team performance within realistic cybersecurity scenarios, while providing the needed metrics to measure cybersecurity competencies quantitatively. While this study achieved strong consensus, like any other research, several limitations were observed, including a relatively small SME panel size (n=20) and the absence of empirical testing with users. Future research will employ hands-on cyber range experiments to measure the FCCI by comparing KSTs measured across human-only and human–GenAI teams. Ultimately, this research advances cybersecurity workforce development by establishing a validated foundation for a quantitative assessment of cybersecurity competencies based on DCWF

necessary for effective collaboration between humans and GenAI in defending against complex and evolving cyber threats.

November 13th at 4:40 PM

Interactive Cybersecurity Lab: Hands-on Cybersecurity Training

Chrstian Soucy, Danai Chasaki

Cybersecurity is a dynamic and essential discipline focused on protecting data and systems from malicious threats. Its scope spans personal devices, industrial systems, medical technologies, financial data, and Personally Identifiable Information (PII). As its integration into society and industry deepens, emerging vulnerabilities demand a skilled workforce capable of securing critical infrastructure. This proposed training introduces foundational cybersecurity concepts through a structured series of puzzles and challenges, organized into progressive stages.

November 13th at 5:00 PM

XR Meets Cybersecurity: A Framework-aligned Immersive Game for AI and Cyber Training

Tonia San Nicolas-Rocca

Cybersecurity education is critical to ensuring the safety of individuals online and the protection of sensitive information and digital resources. Given the growing complexity and frequency of cyber threats, there is a pressing need for accessible and engaging training solutions that can effectively reach users across all educational backgrounds, regardless of their geographic location. This work in progress paper presents an approach to the design of an extended reality (XR) immersive training game that integrates cybersecurity and artificial intelligence (AI) concepts. The proposed system is grounded in the National Initiative for Cybersecurity Education (NICE) Workforce Framework for Cybersecurity and the National Centers of Academic Excellence in Cyber Defense (CAE-CD) and Cyber AI Knowledge Units [19, 20, 40]. By leveraging XR technologies, this training game aims to promote cybersecurity and AI awareness, cultivate workforce readiness, and stimulate broader interest in the cybersecurity field.