

A Defense First Threat Speed Training Model for Teaching AI Era Cybersecurity to Non Specialized Professionals

Sarah Sharif

University of Oklahoma

Mohamad Yazan Sadoun

University of Oklahoma

Arash Ahmadi

University of Oklahoma

Sara Aminpour

University of Oklahoma

Logan Larsh

University of Oklahoma

Danial Ebrahimzadeh

University of Oklahoma

Reza Ahmadvand

University of Oklahoma

Meraj Ahmed

University of Oklahoma

Gourab Datta

University of Oklahoma

Tonmoy Malakar

University of Oklahoma

Jonas Manchester

University of Oklahoma

Sumaia Mishu

University of Oklahoma

Yaser Mike Banad

bana@ou.edu

University of Oklahoma

Keywords: cybersecurity education, professional development, self-efficacy, dual-use safety, artificial intelligence security, post-quantum cryptography

Posted Date: October 14th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10910289/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

A Defense First Threat Speed Training Model for Teaching AI Era Cybersecurity to Non Specialized Professionals

Sarah Sharif* Mohamad Yazan Sadoun* Arash Ahmadi
Sara Aminpour Logan Larsh Danial Ebrahimzadeh
Reza Ahmadvand Meraj Ahmed Gourab Datta
Tonmoy Malakar Jonas Manchester Sumaia Mishu
Yaser Mike Banad†

INQUIRE Lab, School of Electrical and Computer Engineering,
University of Oklahoma, Norman, Oklahoma, USA

Abstract

Cybersecurity threats now evolve faster than most professional training can absorb them: AI-enabled attacks, synthetic identity, post-quantum transition risk, and widening wireless and device-facing attack surfaces. The professionals most exposed are often not security specialists, and conventional training seldom reaches them in time. This paper presents the *Threat-Speed Training Model*, a short, defense-first, misuse-resistant instructional design that introduces the breadth of these fast-moving threats to non-specialized professionals with no prerequisites. The model organizes emerging threats into durable risk categories and teaches each through a repeating attacker’s-logic cycle grounded in experiential-learning theory, paired with a concrete, reusable procedure for keeping dual-use content safe for a general audience. We evaluate the model formatively across four cohorts over two years (two annual offerings of two weekly sessions each). The program reached 116 registrants across sectors including tribal-nation government, K–12 and higher education, healthcare, state and federal agencies, and industry, of whom 56 completed a 40-hour certificate. Using post-workshop evaluations, a retrospective self-efficacy measure, a pre-workshop needs assessment, and an expert-evaluator review, we report participants’ self-reported self-efficacy gains, the program’s cross-sector reach, and its reception. As a formative, self-report study with small per-cohort samples, it contributes design knowledge rather than efficacy evidence. We offer (1) the model, with its curriculum, dual-use safety procedure, and adaptable design principles, and (2) a multi-cohort account of how the model reaches and is received by a workforce that conventional training does not.

*These authors contributed equally to this research.

†Corresponding author: bana@ou.edu.

Keywords: cybersecurity education; professional development; self-efficacy; dual-use safety; artificial intelligence security; post-quantum cryptography

1 Introduction

Cybersecurity education is under pressure from a widening gap between the speed of threat change and the speed of formal training (Figure 1). New attack surfaces now emerge from artificial intelligence systems, synthetic media, connected devices, wireless interfaces, identity platforms, cloud services, and cryptographic transition risks. At the same time, vulnerability discovery, disclosure, and exploitation cycles are accelerating through public vulnerability databases, automated scanning, exploit-sharing channels, bug-bounty programs, and AI-assisted code analysis. As a result, cybersecurity professionals must not only understand foundational concepts, but also learn how to interpret and respond to unfamiliar threats as they emerge.




Figure 1: Conceptual schematic; the axes are illustrative and no measured data are plotted. As AI accelerates threat evolution, the capability that conventional training cycles build falls behind; the model is designed to help non-specialized professionals close the distance.

This pressure is intensified by the rapid development of large language models (LLMs) and AI-assisted software tools. The same capabilities that help defenders review code, summarize alerts, and identify vulnerabilities can also help attackers scale reconnaissance, generate convincing phishing content, adapt exploit code, and move more quickly from vulnerability discovery to attempted compromise [1]. As LLMs become more capable and more deeply integrated into software development, cloud services, security tools, and everyday workflows, the systems that learners use and defend face new forms of compromise faster than traditional training cycles can address.

This acceleration changes what cybersecurity training must provide. Learners still need foundations, but foundations alone are no longer enough for professionals facing current threats. Effective training must connect core concepts to recent tactics, tools, and defensive

decisions. It must help participants understand why a password policy, an identity control, a firewall rule, a patch strategy, or a cryptographic migration matters in realistic operational settings. It must also prepare learners to interpret unfamiliar threats after the workshop ends, because many of the risks they will face will not match examples from older curricula.

Emerging risks intensify this need for practical, current instruction. AI-powered social engineering can increase the scale and credibility of phishing, synthetic identity, and deepfake-enabled deception; phishing alone already drives a large and fast-growing share of attacks, and automated detection remains imperfect across regions and languages [2]. Retrieval-augmented generation (AI assistants that pull in outside documents) and other AI-enabled services introduce new risks around data exposure, prompt manipulation, and trust in generated outputs [3]. Quantum-era security raises long-term concerns about cryptographic readiness and migration planning. Wireless, radio-frequency, NFC, USB, and smart-city technologies expand the boundary between cyber, physical, and everyday devices. These topics are technically diverse, but they share a common educational challenge: learners need enough realism to understand the risk and enough safety structure to avoid turning instruction into misuse guidance.

Short-format professional training can address this challenge when it is designed around practical, current, and defense-first learning. We use threat-speed training to describe a short-format model that updates around fast-moving risk categories, teaches them through safe hands-on activities, and translates them into defensive action. A five-day workshop cannot make participants experts in every emerging threat area. It can, however, expose learners to the latest risk categories, show how those risks appear in real systems, provide controlled hands-on practice, and connect each topic to prevention, detection, response, and responsible decision-making. This format fits professionals who work near cybersecurity problems but cannot commit to a full academic course or deep specialization in every topic.

We present the *Threat-Speed Training Model*, a short, defense-first instructional design for this purpose, and evaluate it formatively across four cohorts over two years (two annual five-day offerings, each run as two weekly sessions; 116 registrants). The curriculum combines emerging-trend briefings, hands-on labs, invited expert talks, open discussions, and games, organized around six durable risk clusters spanning AI-era, cryptographic, network and device, identity, detection, and incident-response threats (Section 3).

The study asks what a short, defense-first model can achieve within its format: awareness and defensive judgment rather than specialist skill. Three research questions organize the work. **RQ1 (design):** how can instruction be designed so that non-specialized professionals across diverse sectors engage safely and productively with fast-evolving, dual-use cyber threats? We answer RQ1 as a design contribution: we propose and detail one such design, with the rationale behind its choices. **RQ2 (response):** how do participants respond, in demand, self-reported self-efficacy, and reception, with any behavior change they volunteer? **RQ3 (relevance):** how well do senior practitioners’ views of workforce priorities align with the curriculum’s coverage? We answer RQ2–RQ3 with a learner lens (a needs assessment, post-workshop evaluations, a before-and-after self-efficacy measure rated on one post-workshop form, and qualitative feedback) and an expert-evaluator review. Throughout, we keep two kinds of number distinct: the program’s *reach* (116 registrants, 56 completers) and its *learning evidence* (per-instrument samples of 13 to 22 respondents); the former establishes who the model serves, the latter speaks to what participants report gaining.

This paper makes two contributions. First, it presents and details the Threat-Speed Training Model: its organizing principles, a worked instance of the attacker-logic teaching cycle (Section 3.1), a curriculum map, and, most useful for adopters, a concrete, repeatable procedure for keeping dual-use content safe for a general audience (Section 5). Second, it provides a formative, multi-cohort account of how the model performs in practice: who it reaches across sectors, how participants receive it, and how its coverage aligns with where senior practitioners see the workforce least prepared. We offer the design for adoption and the artifacts for reuse, and report it as design knowledge from one multi-cohort cycle of real delivery.

2 Background and Related Work

Cybersecurity moves faster than the programs meant to teach it. National workforce frameworks define the cybersecurity workforce and its skill areas in detail [4], and recent meta-analytic work finds that security training reliably improves end-user knowledge and awareness but struggles to change behavior [5]. The open problem is how to keep *non-specialized* professionals current with threats that did not exist a year earlier, without turning every lesson into a specialist course. Our work addresses that gap by integrating five areas: hands-on and short-format training, AI-era threat education, ethical dual-use instruction, expert-panel curriculum validation, and self-report methodology.

Hands-on and short-format security education. Much hands-on security training takes the form of cyber ranges, capture-the-flag exercises, and skills laboratories that build and measure operational technique [6]. These are effective for producing security *practitioners* (typically university students or specialists working toward a technical role), but they assume time, prerequisites, and that trajectory. Broader-audience initiatives exist (the NSA/NSF GenCyber summer academies introduce K–12 students and teachers to cybersecurity [7], and CyberCorps: Scholarship for Service funds a degree-track pipeline into government roles), but they serve students and future specialists, not the working professional already in a non-security role. A short, no-prerequisite workshop for working professionals across many sectors is a different instrument with a different goal: awareness and defensive judgment rather than operational skill.

Teaching emerging, AI-era threats. Where the newest threats are taught, they tend to be taught one at a time. Deepfakes have their own awareness studies and simulations [8], with the underlying technology separately reviewed [9]; post-quantum cryptography, adversarial machine learning, and large-language-model security each have dedicated treatments. Each is valuable, and each is narrow. No program bringing this *breadth* of AI-era threats (synthetic identity, adversarial images, wireless and proximity attacks, post-quantum readiness, and more) into one coherent curriculum for a general professional audience surfaced in our review of the literature; designing and evaluating one is the gap this work addresses.

Safe and ethical instruction. Teaching how attacks work raises an obvious dual-use concern, and the community has set out principles for delivering offensive content responsibly [10]. Those principles are sound but largely conceptual; this paper turns “defense-first” from a stance into a delivered, evaluated, repeatable instructional design (Sections 3 and 5).

Expert validation of curricula. Establishing that a curriculum covers the right mate-

rial is itself a research problem, and Delphi and expert-panel methods are the standard tools, including in cybersecurity education [11, 12]. These studies validate *what to teach*. We apply the idea in a complementary direction: an expert-evaluator panel rates each topic on both its importance and the workforce’s current preparedness, producing a gap the curriculum can be checked against, an approach others can apply.

Measuring perceived change and self-efficacy. The learner-side construct we measure is *self-efficacy*: a person’s belief in their capability to organize and execute the actions a situation demands [13]. Self-efficacy is a recognized outcome in professional-development research, is measurable by self-report, and matters practically because it predicts whether people attempt, persist at, and act on what they have learned, which for a non-specialist audience is precisely the pathway from awareness to defensive action. Where a true pre-test is impractical, the retrospective then-test (rating “before” and “now” on one post-workshop form) is a recognized design we adopt [14, 15].

Each thread is well developed alone; the contribution here is their integration into one delivered, evaluated model: a short, defense-first, misuse-resistant program that teaches this breadth to non-specialized professionals with a concrete procedure for keeping dual-use content safe.

3 The Threat-Speed Training Model

We use *threat-speed training* to name a short-format, defense-first instructional model designed to help non-specialized professionals keep pace with rapidly evolving cyber threats. The model does not aim to make participants experts in any single area. Instead, it exposes them to the current categories of risk, shows how those risks appear in real systems, gives controlled hands-on practice, and connects each topic to prevention, detection, response, and responsible decision-making. Figure 2 shows what the model covers (its risk clusters and the defense-first boundary within which they are taught), and Figure 3 the per-topic teaching cycle.

The model rests on five design principles. Several are individually familiar from good laboratory practice; the model’s novelty is in their integration, above all the dual-use safety procedure of Section 5, which operationalizes them for a general, non-operator audience.

P1. Organize around fast-moving risk categories, not isolated tools. Emerging threats are grouped into durable clusters (AI-enabled and synthetic-identity threats; cryptographic and post-quantum readiness; network, wireless, and device-facing attack surfaces; identity and access control; vulnerability triage and intrusion detection; and incident response) so the curriculum can be refreshed at the speed of the threat landscape without redesigning the program each year.

P2. Teach the attacker’s logic, in service of defense. The workshop goes beyond showing *that* an attack works to explaining the attacker’s reasoning (why a particular target is chosen and how the exploit actually succeeds), because understanding that logic is what lets a non-specialist recognize an unfamiliar variant later. Every such demonstration is paired with an explicit defensive takeaway (a control, a detection signal, or a policy) as worked through in Section 3.1; Table 1 summarizes these at the daily level. The goal is awareness and prevention, not operational capability.




Figure 2: The Threat-Speed Training Model at a glance: six durable risk clusters (principle P1), taught for prevention, detection, response, and responsible practice, all within a defense-first, misuse-resistant boundary, and delivered as a five-day hands-on workshop.

P3. Make hands-on activities safe by construction. All practical work runs in controlled, isolated environments on the program’s own devices. There are no real third-party targets, no live credentials, and no operational exploitation (Section 5).

P4. Pair instruction with practitioner perspective. Expert-led sessions are complemented by invited talks from working professionals across the public, private, state, and federal sectors, grounding each topic in real workplace consequences.

P5. Lower the barrier to entry. The program requires no prerequisites and is built for a mixed audience, using plain-language framing, daily reflection, and interactive games to support retention.

Within each session these principles are realized as a five-step micro-cycle: *concept* → *safe demonstration* → *guided hands-on* → *defensive framing* → *reflection*. Every topic runs the concept, defensive-framing, and reflection steps; the safe-demonstration and guided-hands-on steps are both included where a technique meets the safety rules of Section 5, and replaced by an instructor-led demonstration where it does not (the Supplementary Information gives the resulting module formats). The cycle moves a learner from “I have heard of this” to “I have seen how it works and know what to do about it” without crossing into misuse guidance.

This cycle draws on experiential-learning theory [16], with each lab *scaffolded* (the instructor withdraws structure as competence grows [17]) to keep a no-prerequisite, mixed-experience audience productive. We illustrate the full cycle on one topic in Section 3.1.




Figure 3: The Threat-Speed Training Model. Each emerging threat is taught through this five-step cycle, with the demonstration and hands-on steps adapted per topic to the safety rules of Section 5: *concept* (what the threat is), *safe demonstration* (seeing it in action without risk), *guided hands-on* (trying it in a controlled lab), *defensive framing* (how to detect and prevent it), and *reflection* (open discussion that connects it to the learner’s own work), turning unfamiliar threats into informed, defensive recognition. Across every step, the model emphasizes the attacker’s logic (how an exploit works), not just that an attack succeeds. The cycle repeats for every topic across the five days.

3.1 The cycle, worked: an adversarial-examples lab

To make the model concrete, we trace one topic, adversarial examples against image classifiers, through the full cycle.

Concept. Participants first meet the idea that machine-learning models can be fooled by inputs that look unchanged to a person, and why this matters now that classifiers sit inside security pipelines (malware triage, content moderation, biometric checks).

Attacker’s logic (P2). Rather than stopping at “the model is wrong,” the session explains *why* the attack works: the attacker nudges the image by a tiny amount, chosen to push the model toward the wrong answer, while the change stays too small for a person to notice (the fast gradient sign method). Grasping this lets a non-specialist recognize the wider family of evasion attacks later, not just this instance.

Safe demonstration. The instructor shows a stock image that MobileNet classifies correctly, applies a change too small for the eye to see on the program’s own machine, and the same model now reports a confident but wrong label.

Guided hands-on. In a scaffolded notebook (pre-loaded model and helper code, run in the isolated lab environment with no third-party systems), participants adjust the size of the change and watch the predicted label and confidence change. The lab occupies a session on the program’s own devices.

Defensive framing. The takeaway is operational: AI components need their own security

testing, and defenses include adversarial training, input- and confidence-range validation, and keeping a human in the loop for high-stakes automated decisions.

Reflection. The session closes with open discussion connecting the lab to participants’ own systems: where a classifier or other AI component sits in their organization’s workflow, and who would own the security testing this lab just demonstrated the need for.

This is the level of detail an adopting instructor needs to reproduce one cycle; the day-by-day module map appears in the Supplementary Information.

4 Workshop Curriculum and Daily Structure

The workshop runs five consecutive days, each from morning to late afternoon, offered each year as two interchangeable weekly sessions to widen access. Every day follows the same rhythm: check-in and breakfast for informal networking, a sequence of expert-led sessions interleaving the micro-cycle of Section 3, an invited practitioner talk over lunch, and a closing block of interactive games and guided reflection that consolidates the day’s defensive takeaways.

Across the five days the curriculum traverses the risk-category clusters of principle P1, from foundations and media trust to defensive tooling, identity, and social engineering. Table 1 maps each day to a representative hands-on activity and its defensive takeaway; the day-by-day module map appears in the Supplementary Information. The invited talks, delivered by a retired federal law-enforcement agent, state and industry security leaders, and senior officials, recur across the offerings and consistently rank among the components participants found most valuable (Section 7).

5 Safety, Ethics, and Misuse-Resistance

Teaching how attacks work can become misuse guidance if exercises are not bounded. Because the audience comprises general professionals rather than security operators, the program is explicitly *misuse-resistant*: it is designed so that a participant leaves better able to recognize and prevent an attack, not to carry one out [10].

Table 2 lists the dual-use controls applied to every hands-on activity; they put principles P2 and P3 into practice. Two rules governed whether a topic became a full hands-on lab or a demonstration. A technique was given a hands-on lab only when (i) it could run entirely self-contained on program hardware, so the exercise never touches a real-world target, and (ii) the awareness goal was reachable below the depth needed to weaponize it. Adversarial examples (Section 3.1) and the machine-learning-defended SQL-injection lab meet both, because the “target” is a system the program ships. Techniques that fail the first rule were reframed toward detection and prevention rather than dropped: the evil-twin Wi-Fi and USB-injection topics, for example, are taught as bounded demonstrations of how the attack deceives users and the controls that stop it, not as payload authoring. This preserves coverage of the threat landscape without lowering the safety floor.

The contribution is not the individual controls (four are standard hands-on-lab hygiene) but their integration for a general, non-operator audience: the combination is what makes

Table 1: Five-day curriculum map. Each activity is paired with a defensive takeaway; hands-on and demonstration activities run under the safety controls of Table 2 (the Supplementary Information gives each activity’s module format).

Day	Focus theme	Representative activity	Defensive takeaway
1	Foundations & media trust	Steganography / hidden-data detection; password-manager and multi-factor authentication setup	Verify the authenticity of digital media; harden personal and account security
2	Web & AI-application attacks	SQL-injection exploit-and-defend on a vulnerable login versus a machine-learning-defended version; retrieval-augmented-generation (RAG) attack walkthrough; vulnerability triage and patch exercise	Defend web applications and AI assistants; prioritize patching by risk
3	AI-era & physical-layer threats	Adversarial-image (FGSM) lab fooling an image classifier; post-quantum key-exchange game; NFC / USB “digital pickpocket” station	Recognize AI-model, quantum, and proximity-device risks and their mitigations
4	Network defense & cryptography	DDoS, intrusion-detection, and incident-response session; RF/IR “invisible world” demonstration; cryptography foundations and post-quantum migration	Detect intrusions and respond; reason about encryption and cryptographic readiness
5	Tools, identity & social engineering	Network-security-tools lab; synthetic-identity and AI-driven social-engineering session; evil-twin Wi-Fi defense; identity-and-access management with a hardware root of trust	Apply defensive tooling; resist social engineering; manage identity and trust

each offensive technique teachable to non-specialists at the depth needed for recognition, bounded below the depth needed for operational capability, the distinction that prior ethical-instruction guidance leaves conceptual [10]. Figure 4 shows how each dual-use topic is taught above this misuse boundary. As learner-side evidence that this delivery was received as intended, participants on the 2026 awareness form agreed the workshop was “taught safely and responsibly” (mean 4.54/5, $n = 13$).

6 Study Design and Data Sources

6.1 Design

We study the model through a multi-cohort formative program evaluation: the program is a designed intervention, delivered as two annual offerings and evaluated formatively against its goals rather than tested for a generalizable effect. The work draws on the design-based tradition in education research [18], but we claim formative evaluation rather than iterative design-based research: the curriculum evolved between offerings, informed by participant feedback and expanding topic coverage, but this refinement was not run as a structured design-based-research cycle with a documented rationale and measured effect for each change.

Table 2: Dual-use controls applied to every hands-on activity.

Control	Implementation
Isolated environment	Program-owned devices; no production or third-party systems
No live targets	Self-contained labs; attacks never directed at real services
Sanitized / simulated data	No real credentials, personal data, or operational artifacts
Defense-paired demonstrations	Every offensive technique mapped to a control or detection signal
Bounded depth	Coverage sufficient for awareness; weaponization detail withheld
Reflection & responsible-use framing	Daily closing reinforces ethical and preventive intent

Table 3: Research questions mapped to data sources and reporting locations.

Research question	Data / basis	Reported in
RQ1 (design): how can instruction be designed for safe, productive engagement with dual-use threats?	The model, its principles, the worked cycle, and the dual-use safety procedure (a design contribution, not a comparative test)	Sections 3–5
RQ2 (response): how do participants respond?	Needs assessments ($N=137$, $N=13$); post-evaluations ($N=22$ /year); retrospective self-efficacy block ($n=13$); open-ended feedback	Sections 7.1–7.5
RQ3 (relevance): do expert evaluators’ views of workforce priorities align with the curriculum?	Expert-evaluator panel importance-vs-preparedness ratings ($n=4$)	Section 7.6

The evaluation instruments were extended in the same spirit, with the 2026 offering adding the retrospective self-efficacy block and the awareness-form agreement items to measure what the 2025 evaluation could not. The evaluation uses two lenses. The *expert-evaluator review* asks senior security practitioners which emerging threats matter most for the workforce and where it is least prepared, as a check on the relevance of the curriculum’s coverage [11, 12]. The *learner lens* asks participants how the program affected their awareness and self-efficacy and how they received it, with open-ended room to report any change in their own behavior. Together the two lenses inform whether the design targets the right threats and whether learners report meaningful change. Table 3 maps each research question to the data that answer it. Because the same team designed, delivered, and evaluated the program, we frame the findings as design knowledge rather than independent efficacy evidence (Section 9).

6.2 Setting and participants

The workshop was delivered at a single university¹ as two annual five-day offerings over two years (2025 and 2026), each run as two weekly sessions to widen access, giving four cohorts

¹Program and registration details: <https://inquirelab.ai/cybersecurity/>.




Figure 4: How each dual-use topic is taught above the misuse boundary: an emerging threat moves to an instructor-led safe demonstration, to a sandboxed hands-on lab, to a defensive decision and reflection, building defensive understanding without crossing into operational misuse.

across the two years. Across the offerings, 116 professionals registered and 56 completed the program (a 40-hour certificate). Not every completer returned every instrument, so Table 4 reports each instrument’s response count alongside enrollment, keeping response coverage explicit. Registrants spanned a wide range of sectors with no security prerequisites: tribal-nation government, healthcare, K–12 and higher education, public agencies, and industry (Section 7.1).

6.3 Data sources and instruments

Table 4 inventories the data. Three families of instrument were used. All instruments were developed in-house for program evaluation (the Supplementary Information summarizes their items); they carry face validity for their descriptive purpose but no external validation or reliability evidence, which is one reason the analysis stays descriptive (Section 9). Registration captured affiliation and role only, so participant characteristics beyond sector and prior experience were not collected.

Pre-workshop needs assessment. Prospective participants rated how strongly they wanted each of a broad set of topics emphasized (1–5): an open 2025 pre-registration survey and a shorter 2026 enrolled-participant version (Table 4). It is a needs/demand instrument, not a knowledge test; it characterizes *demand*, not learning.

Post-workshop evaluation. Participants rated the workshop on seven five-point items, selected the components or topics they found most valuable, and answered open-ended questions. The 2026 Week-1 form asked for most-valuable *topics* and the Week-2 form for most-valuable *components*, reported separately.

Retrospective self-efficacy self-assessment. The 2026 awareness form added a within-person retrospective “then/now” block: for each of ten threat topics (the same taxonomy

Table 4: Data sources. Enrollment and survey-response counts are reported separately; instruments differ by year.

Source	Year/session	<i>N</i>	Measures
Registrations	2025+2026	116	Affiliation, role
Completed program	2025+2026	56	40 training hours
Needs assessment [†]	2025	137	Topic-priority demand
Needs assessment	2026	13	Topic-priority demand + role
Post-evaluation	2025	22	Satisfaction, value, open-ended
Post-evaluation	2026	22	Satisfaction, value, open-ended
Before/now self-efficacy	2026 (W1+W2)	13	Within-person confidence change
Expert-evaluator panel	2026	4	Importance vs. knowledge gap

[†]The 2025 needs assessment was an open pre-registration interest survey reaching a broader prospective audience; its 137 respondents are not a subset of the 116 registrations. The 2026 version was completed by enrolled participants.

used by the expert-evaluator panel), participants rated their confidence *before* the workshop and *now*, on one form. We interpret these ratings as self-efficacy [13], and the retrospective then-test is a recognized design for self-reported change [14, 15], which we adopt as supporting evidence.

Expert-evaluator panel. A panel of senior security practitioners, serving as expert evaluators, rated each of the same ten topics on how important it is and how much the workforce knows, yielding a knowledge gap. They evaluated the curriculum’s content and took no part in designing or preparing the curriculum. We report their review in Section 7.6.

6.4 Analysis

Given small per-instrument samples, analysis is descriptive. For satisfaction and value we report response distributions per year (Table 5; full distributions in Supplementary Table S1). For the retrospective block we report, per topic, the mean “before” rating, the mean “now” rating, the mean within-person change, and the proportion of respondents rising at least one point; the figures also show every individual response so the spread behind each mean is visible. At these sample sizes a single respondent shifts a topic mean by roughly 0.08 ($n = 13$) and a panel gap by 0.25 ($n = 4$), so all per-topic values should be read as indicative descriptions of this group, not stable estimates. For the expert-evaluator panel we report rankings and the importance–preparedness gap, treating it as expert judgment rather than a generalizable survey. Open-ended responses were analyzed using thematic analysis [19]: responses were coded inductively, codes were grouped into the themes of Table 6, and quotations illustrate themes the quantitative results also support.

6.5 Ethics

The study was reviewed by the University of Oklahoma Institutional Review Board and determined to be exempt from IRB review under Exempt Category 4, 45 CFR 46.104(d)(4) (IRB #20222, approved August 12, 2026), because it involved only secondary research use of information originally collected for voluntary program evaluation, recorded in such a



Figure 5: Registrants by sector across the two years. Sectors were coded by keyword matching on each registrant’s listed organization; 105 of the 116 registrants named an organization. No prior security background was required.

manner that the identity of participants could not readily be ascertained. The study was carried out in accordance with the institutional guidelines of the University of Oklahoma Institutional Review Board, the U.S. Federal Policy for the Protection of Human Subjects (45 CFR part 46), and the ethical principles of the Declaration of Helsinki. The surveys were administered as voluntary program evaluation with freely given, informed consent. Providing a name was optional on all surveys, responses were analyzed and reported only in aggregate, and the participant excerpts in Table 6 are unattributed. The expert evaluators gave their assessments with knowledge of the study; the quotations attributed to them by role and sector are used with their explicit permission, and one who declined attribution is not quoted.

7 Results

7.1 Reach and participant diversity

A central aim of the model is to reach professionals who work near cybersecurity but are not specialists, and the registration record bears this out: the 105 of 116 registrants who named an employer span higher education, state and federal government (among them a state homeland-security office and a military installation), tribal-nation government, K–12 school districts, healthcare, and industry (Figure 5). Of the 2026 needs-assessment respondents who reported an experience level (11 of 13), nine had three years or fewer (four of them less than one year) and two had eight or more, consistent with the no-prerequisite design.

7.2 Self-assessed self-efficacy change

On the 2026 awareness form’s retrospective block ($n = 13$, spanning Week 1 and Week 2), mean self-assessed confidence rose on every one of the ten threat topics (Figure 6). Mean gains ranged from +0.62 to +2.23 on a five-point scale, and on every topic except account security, at least nine of the thirteen participants rose by at least one point, all thirteen for AI-powered phishing. With thirteen respondents these means are indicative rather than precise (Section 6.4); the individual responses plotted behind each mean in Figure 6 show the spread directly. The largest gain was for the quantum threat to encryption, which also started from the lowest base ($1.69 \rightarrow 3.92$); the smallest was for account security, which began near the top of the scale (4.23) and left little room to move (only five of the thirteen rose there). The largest gains fall where prior confidence was lowest, consistent with the program introducing genuinely unfamiliar material.

A separate item agrees, from a different instrument: on the 22-respondent post-evaluation form (not the 13-respondent retrospective block), participants selected the topics they already knew. The largest gain (the quantum threat to encryption) was among the least familiar going in (2 of 22 reported prior knowledge), and the smallest gain (account security) the most familiar (password management, 7 of 22).

Two further items track the model’s durable goal more directly than topic confidence: participants’ “realistic, practical sense of how these attacks actually happen” rose +1.2 ($3.08 \rightarrow 4.31$), and on a post-workshop agreement item they endorsed knowing where to keep learning about new threats (4.46/5), the transferable recognition the model aims for beyond any single topic.

7.3 Reception across the evaluation

The post-workshop form carried seven five-point items each year: overall satisfaction, content relevance, clarity, presenter effectiveness, interactive learning, hands-on value, and likelihood to recommend (Table 5). On every item in both years, 86 to 100% of participants chose one of the two most positive options (favorable). Across both years, hands-on activities and invited expert speakers were the components participants most often named as valuable, consistent with the model’s hands-on, practitioner-led emphasis. On the 2026 awareness form the same emphasis appeared as direct agreement: participants agreed the hands-on activities aided understanding (4.38/5), the speakers connected the material to real work (4.38/5), and the five-day format worked well for learning current threats (4.54/5).

7.4 Demand for the curriculum’s topics

The pre-workshop needs assessments (an open 2025 interest survey, $N = 137$, and a shorter 2026 enrolled-participant survey, $N = 13$, reported separately rather than pooled) both show strong, broad demand for current and emerging-threat topics, with AI-powered attacks and defenses among the highest priorities. Because respondents self-selected into a cybersecurity interest survey, we read this as demand context, confirming that the curriculum covers topics this audience wants, not as independent evidence that the curriculum is well-targeted, and not as evidence of efficacy.




Figure 6: Retrospective self-assessed confidence, before vs. after the workshop ($n = 13$, 2026 Week 1 + Week 2), sorted by gain. Open dots are the mean “before” rating and filled dots the mean “now”; faint dots behind each show the individual respondents, so the small-sample spread is visible. The right-hand column is the mean within-person gain. Reported as perceived change (self-efficacy), not objective gain; with $n = 13$ the means are indicative rather than precise.

7.5 Qualitative feedback

Open-ended responses were coded thematically; Table 6 summarizes the recurring themes with representative excerpts. The strengths echo the quantitative picture: participants singled out the distinctive hands-on activities, valued the outside professional perspective, and connected the material to their own work, with several asking to return. Asked to name one thing they would now watch for, participants gave specific, on-topic answers spanning the curriculum: researching NFC and tap-payment risks, scrutinizing email attachments more closely, and preparing for quantum-era cryptographic change. These self-reported shifts show the workshop produced concrete defensive intentions, though they do not measure downstream behavior. The improvement requests were specific and constructive (more forensics and tabletop incident-response practice, deeper network-defense content, and more visual, less text-heavy delivery).

Table 5: Post-workshop reception: the percentage of participants choosing one of the two most positive options (“favorable”) on each five-point item, of $N = 22$ valid responses per year. Reception was strongly positive on every item in both years; because responses were voluntary, these figures describe the respondents, who may overrepresent the program’s more engaged participants (Section 9). Reported per year, never pooled; full per-item response distributions appear in Supplementary Table S1.

Item (5-point)	2025	2026
Overall satisfaction	95%	100%
Content relevance	86%	86%
Clarity of material	91%	95%
Presenter effectiveness	86%	91%
Interactive learning	95%	95%
Hands-on value	86%	91%
Would recommend	91%	95%

7.6 Expert-evaluator panel

An expert-evaluator panel rated each of the ten threat topics on its importance to the workforce and the workforce’s current preparedness. The four evaluators are senior security practitioners with backgrounds including enterprise security leadership (a former Fortune-500 chief information security officer), law enforcement (a serving police chief and former military instructor), and offensive security (a penetration-testing-firm executive). They rated topic importance and preparedness only and took no part in designing or preparing the curriculum. We present their ratings as an expert evaluation of the curriculum’s relevance.

The panel unanimously rated keeping professionals current as “critical” and endorsed the defense-first approach. One, a former chief information security officer, captured the rationale: “A short, hands-on, defense-first workshop directly addresses the gap between exposure and readiness. It helps non-specialists practice recognition, verification, safe reporting, and defensive habits without turning the session into offensive training.” Figure 7 shows the resulting knowledge gap: the largest, where the workforce is least ready, fall in the AI-era topics: AI-powered phishing (+2.0), deepfakes (+1.8), and fooling AI image recognition (+1.8). That AI-era cluster holds the evaluators’ largest preparedness gaps and several of the participants’ largest self-efficacy gains (Figure 6); the single largest individual gain was for the quantum threat to encryption, a different cluster. Section 9 explains why we read this partial overlap as consistency rather than triangulation.

8 Discussion

Across both lenses, the three research questions are answered within the study’s formative scope: RQ1 (design) by the model and its dual-use safety procedure, presented as a proposed and detailed design with its rationale; RQ2 (response) by strong demand and reception alongside self-efficacy gains on all ten topics with concrete defensive takeaways; and RQ3 (relevance) by the alignment of expert evaluators and learners on the AI-era cluster. On RQ3, we read that alignment as consistency rather than strong corroboration: both lenses

Table 6: Themes from open-ended responses. Quotation marks indicate verbatim excerpts; unquoted entries are coded summaries of a theme’s typical content.

Theme	Representative excerpt
<i>Strengths</i>	
Intended behavior change	“I question attachments even more now”; “I understand the very real risks of quantum computing and how to reduce their effect on my organization.”
Hands-on activities valued	“Got some new tools for my tool belt.”
Accessible to non-specialists	“Great job explaining topics from a baseline of little to no prior knowledge.”
Eye-opening on emerging threats	“Technology is advancing much faster than I had originally anticipated.”
Knowledgeable, well-organized delivery	“The trainers were very knowledgeable, and the workshop was very well organized.”
Outside professional perspective	Invited speakers repeatedly named among the most valuable parts.
Strong recommendation & repeat value	“An outstanding experience... I will be recommending this workshop to others”; “my second year attending and I still found it beneficial.”
Relevance / intent to return	“Can I attend again next year?”
<i>Improvement requests</i>	
More forensics / capture-the-flag	Hands-on forensic-tool exercises requested.
Tabletop incident response	“Adding a tabletop exercise for incident response would be beneficial.”
Deeper network-defense content	Firewall, VLAN, and defensive-design depth requested.

are structurally drawn toward the newest topics, which start least familiar to learners and are judged least covered across the workforce, so the agreement partly reflects that shared structure. Consistency is still informative, since both lenses could have pointed elsewhere, but it falls short of triangulation.

The learner-lens results form a consistent pattern. Participants reported gains across the full breadth of topics, and the gains were largest on the topics where prior confidence started lowest: quantum readiness above all, with AI-powered phishing, fooling AI image recognition, and steganography close behind. These are topics that single-subject courses rarely cover together and that a general professional is unlikely to encounter elsewhere. The pattern, within the study’s self-report limits, is consistent with the model’s central aim: to move non-specialized professionals from unawareness to informed, defensive recognition across a wide and shifting threat surface. This awareness-and-recognition target is also where the evidence base is strongest: a 69-study meta-analysis finds that security training reliably improves knowledge and awareness but rarely demonstrates behavior change [5]. A program that moves self-efficacy across a broad threat surface is therefore targeting what such training can most credibly achieve, and we read the self-reported behavioral takeaways (Section 7.5) as a modest signal toward the harder outcome, not a behavior-change claim.

The needs assessments show that the curriculum’s topics match what this self-selected audience says it wants, useful demand context for adopters even if not evidence of targeting. The model is built to be adopted: organizing content around durable risk *categories* (P1)



Figure 7: Expert-evaluator panel: how important each topic is vs. how much the workforce knows, sorted by gap. Each expert evaluator rated every topic on its importance to the workforce and the current workforce’s preparedness for it, from 1 (very low) to 5 (very high); filled dots are mean importance, open dots are mean preparedness (read here as how much the workforce knows), and the knowledge gap is importance minus preparedness. Faint dots show the individual evaluators behind each mean; with $n = 4$, one rater moves a gap by 0.25. The AI-era topics top the list; participants’ largest self-efficacy gains fall in the same cluster.

lets other institutions refresh material at threat speed without redesigning the program, and the defense-first controls (Section 5) make it adoptable in settings wary of teaching attack techniques. The panel reinforced the adoption case: one evaluator recommended a modular, role-based packaging, with awareness-, practitioner-, and authorized-lab-level versions of each topic plus instructor notes, safety boundaries, and assessment rubrics, so the same threat themes transfer across an organization without exposing every audience to the same operational depth. Concretely, an adopter needs three things, in order: the risk-category map and the worked cycle (Sections 3–3.1) to structure content; program-owned lab hardware and the dual-use safety procedure (Section 5) to keep activities misuse-resistant; and instructors able to demonstrate each technique at awareness depth. We present this as design knowledge from one multi-cohort cycle of real delivery; external replication at a second site is the next validation step, and the same importance-versus-preparedness comparison can help an adopter check coverage against a local workforce gap.

9 Limitations

These findings are formative design knowledge from a single program, and several features of the design bound what they can support.

First, the study includes no objective measure of learning: every learner-side outcome is self-reported. We interpret the confidence ratings as self-efficacy [13], a recognized but perceptual construct, and the retrospective then-test mitigates but does not remove response-shift bias [14, 15]; self-report from volunteers can also carry social-desirability effects. We therefore claim perceived learning, not objective knowledge or skill, and no downstream behavior or security outcome. The volunteered behavioral takeaways (Section 7.5) signal intent toward the harder outcome that the broader evidence base finds training rarely achieves [5].

Second, the samples are small and uneven, which is why the analysis is descriptive. The retrospective self-efficacy instrument has 13 respondents and the expert-evaluator panel four; at these sizes a single respondent shifts a topic mean by roughly 0.08 and a panel gap by 0.25, so the per-topic values are indicative descriptions of these groups rather than stable estimates, and the figures plot every individual response so the spread is visible. Reach and learning are also distinct constructs in this study: the cross-sector reach (116 registrants, 56 completers) is well supported by the registration record, while the self-efficacy evidence rests on one instrument in one year.

Third, completion and response were partial. Of 116 registrants, 56 completed the certificate, and 22 participants per year returned the post-evaluation. The evaluated sample may therefore overrepresent the program’s more engaged participants, and the reception and self-efficacy measures should be read with that selection in mind; the evaluation data do not let us characterize non-completers.

Fourth, the expert-evaluator panel is small: four raters, whose ratings we treat as expert judgment rather than a generalizable survey. Because both lenses are structurally drawn toward the newest topics, the agreement between the panel’s largest perceived workforce gaps and participants’ largest self-efficacy gains should be read as consistency, not strong triangulation. A larger external panel remains a priority for future work.

Fifth, the curriculum itself evolved between the two years, expanding topic coverage in response to participant feedback, so the 2025 and 2026 offerings are not identical programs; the self-efficacy instrument also ran only in the 2026 cohorts, so conclusions about learning rest on a narrower slice of the program, and on its more recent form, than the reach data.

Finally, all data come from a single university, and the evaluation was conducted by the same team that designed and delivered the program. We state this positionality plainly and weight the paper’s claims toward design knowledge accordingly; external replication at a second site, with an objective knowledge measure and an independent panel, is the natural next step.

10 Conclusion

Cybersecurity training has to move at the speed of the threats it teaches. We presented and evaluated the Threat-Speed Training Model: a short, defense-first, misuse-resistant approach that teaches the breadth of AI-era threats (from synthetic identity and adversarial

images to wireless, post-quantum, and beyond) to non-specialized professionals across many sectors. Across the two annual offerings, participants reported strong value in the hands-on and expert-led components and broad self-assessed gains in self-efficacy, and the needs assessments show demand for the curriculum’s topics among the professionals it serves. By equipping non-specialized professionals to recognize and prevent these attacks, the program aims at a security benefit for their workplaces and the communities they serve, an outcome this formative study motivates but does not measure. Because the model is short, prerequisite-free, and built from reusable parts, other universities and institutions can stand up similar training; we offer the model for adoption. A relevance review by senior expert evaluators points to the same AI-era cluster the curriculum emphasizes; a larger expert panel and an objective knowledge measure are the next steps to test that alignment and the learning the program aims to produce. Teaching emerging threats safely, at threat speed, is both necessary and achievable; this work contributes a model for doing it, and a two-year, multi-cohort account of real delivery showing that it reaches a workforce conventional training does not.

Acknowledgments

We thank our invited speakers (Agent Harris Wesley, retired, FBI; Dr. Curtis Coleman; Dr. Stacey White; Joan Stanolis and Jordan Caldwell, Critical Fault; Mr. Kris Larsh; and MG Tommy Mancino) for their contributions to the program.

Declarations

Funding. This workshop was supported by the Federal Emergency Management Agency (FEMA) and the Oklahoma Office of Homeland Security.

Competing interests. The authors have no competing interests to declare that are relevant to the content of this article. The expert evaluators were invited speakers in the program; they had no role in the curriculum design and no financial interest in the evaluation.

Ethics approval. The study was reviewed by the University of Oklahoma Institutional Review Board and determined to be exempt from IRB review under Exempt Category 4, 45 CFR 46.104(d)(4) (IRB #20222, approved August 12, 2026), because it involved only secondary research use of information originally collected for voluntary program evaluation, recorded in such a manner that the identity of participants could not readily be ascertained. The study was carried out in accordance with the institutional guidelines of the University of Oklahoma Institutional Review Board, the U.S. Federal Policy for the Protection of Human Subjects (45 CFR part 46), and the ethical principles of the Declaration of Helsinki.

Data availability. The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

Materials availability. Summaries of the survey instruments and the day-by-day curriculum module map appear in the Supplementary Information. Additional program information, including current schedules and registration materials, is available at the program website (<https://inquirelab.ai/cybersecurity/>).

Consent to participate. Freely given, informed consent to participate in the study was obtained from all participants. The surveys were administered as voluntary program evaluation; providing a name was optional, and all responses were analyzed and reported only in aggregate. The expert evaluators gave their assessments with knowledge of the study, and quotations attributed to them by role and sector are used with their explicit permission.

Consent to publish. Not applicable. The manuscript contains no individual person’s identifying data, images, or personal details. Participant survey responses are reported only in aggregate without identifying information, and the expert evaluators quoted by role and sector consented to the use of their quotations.

Figures and images. All figures in the manuscript are original and were created by the authors; no third-party images are included, and no permissions or credits are required.

Author contributions. All authors contributed to the design, delivery, and evaluation of the workshop program.

References

- [1] Hanxiang Xu, Shenao Wang, Ningke Li, Kailong Wang, Yanjie Zhao, Kai Chen, Ting Yu, Yang Liu, and Haoyu Wang. Large language models for cyber security: A systematic literature review. *ACM Transactions on Software Engineering and Methodology*, 2025. doi: 10.1145/3769676. URL <https://doi.org/10.1145/3769676>.
- [2] Ying Yuan, Giovanni Apruzzese, and Mauro Conti. Beyond the West: Revealing and bridging the gap between Western and Chinese phishing website detection. *Computers & Security*, 148:104115, 2025. ISSN 0167-4048. doi: 10.1016/j.cose.2024.104115. URL <https://doi.org/10.1016/j.cose.2024.104115>.
- [3] Shenglai Zeng, Jiankun Zhang, Pengfei He, Yue Xing, Yiding Liu, Han Xu, Jie Ren, Shuaiqiang Wang, Dawei Yin, Yi Chang, and Jiliang Tang. The good and the bad: Exploring privacy issues in retrieval-augmented generation (RAG). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 4505–4524, 2024. doi: 10.18653/v1/2024.findings-acl.267. URL <https://doi.org/10.18653/v1/2024.findings-acl.267>.
- [4] Rodney Petersen, Danielle Santos, Matthew C. Smith, Karen A. Wetzel, and Greg Witte. Workforce framework for cybersecurity (NICE framework). Technical Report NIST SP 800-181r1, National Institute of Standards and Technology, November 2020. URL <https://doi.org/10.6028/nist.sp.800-181r1>.

- [5] Julia Prümmer, Tommy van Steen, and Bibi van den Berg. Assessing the effect of cybersecurity training on end-users: A meta-analysis. *Computers & Security*, 150:104206, March 2025. ISSN 0167-4048. doi: 10.1016/j.cose.2024.104206. URL <https://doi.org/10.1016/j.cose.2024.104206>.
- [6] Tiina Schafeitel-Tähtinen and Willi Lazarov. Teaching and learning cybersecurity using capture the flag: Effectiveness comparison between university students in Finland and Czechia. *Computer Applications in Engineering Education*, 33(5), 2025. ISSN 1099-0542. doi: 10.1002/cae.70082. URL <https://doi.org/10.1002/cae.70082>.
- [7] Bryson R. Payne, Tamirat Abegaz, and Keith Antonia. Planning and implementing a successful NSA-NSF GenCyber summer cyber academy. *Journal of Cybersecurity Education, Research and Practice*, 2016(2), 2016. ISSN 2472-2707. doi: 10.62915/2472-2707.1015. URL <https://doi.org/10.62915/2472-2707.1015>.
- [8] Hana Oh, Eunbi Lee, Seungju Shin, Keyeun Lee, Miran Pyun, and Hajin Lim. Deep-Aware: Using experiential deepfake simulations to enhance cybersecurity awareness in older adults. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, CHI '26, page 1–23. ACM, April 2026. doi: 10.1145/3772318.3791492. URL <https://doi.org/10.1145/3772318.3791492>.
- [9] Mika Westerlund. The emergence of deepfake technology: A review. *Technology Innovation Management Review*, 9(11):39–52, November 2019. ISSN 1927-0321. doi: 10.22215/timreview/1282. URL <https://doi.org/10.22215/timreview/1282>.
- [10] Fredrik Heiding and Robert Lagerström. *Ethical Principles for Designing Responsible Offensive Cyber Security Training*, page 21–39. Springer International Publishing, 2021. ISBN 9783030724658. doi: 10.1007/978-3-030-72465-8_2. URL https://doi.org/10.1007/978-3-030-72465-8_2.
- [11] Nabin Chowdhury, Sokratis Katsikas, and Vasileios Gkioulos. Modeling effective cybersecurity training frameworks: A Delphi method-based study. *Computers & Security*, 113:102551, February 2022. ISSN 0167-4048. doi: 10.1016/j.cose.2021.102551. URL <https://doi.org/10.1016/j.cose.2021.102551>.
- [12] Geet Parekh, David DeLatta, Geoffrey L. Herman, Linda Oliva, Dhananjay Phatak, Travis Scheponik, and Alan T. Sherman. Identifying core concepts of cybersecurity: Results of two Delphi processes. *IEEE Transactions on Education*, 61(1):11–20, February 2018. ISSN 1557-9638. doi: 10.1109/te.2017.2715174. URL <https://doi.org/10.1109/te.2017.2715174>.
- [13] Albert Bandura. Self-efficacy: Toward a unifying theory of behavioral change. *Psychological Review*, 84(2):191–215, 1977. ISSN 0033-295X. doi: 10.1037/0033-295X.84.2.191. URL <https://doi.org/10.1037/0033-295X.84.2.191>.
- [14] Frederick R. Rohs, Christine A. Langone, and Rhonda K. Coleman. Response shift bias: A problem in evaluating nutrition training using self-report measures. *Journal*

- of Nutrition Education*, 33(3):165–170, May 2001. ISSN 0022-3182. doi: 10.1016/s1499-4046(06)60187-5. URL [https://doi.org/10.1016/s1499-4046\(06\)60187-5](https://doi.org/10.1016/s1499-4046(06)60187-5).
- [15] Clara C. Pratt, William M. McGuigan, and Aphra R. Katzev. Measuring program outcomes: Using retrospective pretest methodology. *American Journal of Evaluation*, 21(3):341–349, September 2000. ISSN 1557-0878. doi: 10.1177/109821400002100305. URL <https://doi.org/10.1177/109821400002100305>.
 - [16] Alice Y. Kolb and David A. Kolb. Learning styles and learning spaces: Enhancing experiential learning in higher education. *Academy of Management Learning & Education*, 4(2):193–212, 2005. ISSN 1944-9585. doi: 10.5465/amle.2005.17268566. URL <https://doi.org/10.5465/amle.2005.17268566>.
 - [17] David Wood, Jerome S. Bruner, and Gail Ross. The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry*, 17(2):89–100, 1976. ISSN 1469-7610. doi: 10.1111/j.1469-7610.1976.tb00381.x. URL <https://doi.org/10.1111/j.1469-7610.1976.tb00381.x>.
 - [18] Terry Anderson and Julie Shattuck. Design-based research: A decade of progress in education research? *Educational Researcher*, 41(1):16–25, 2012. ISSN 1935-102X. doi: 10.3102/0013189X11428813. URL <https://doi.org/10.3102/0013189X11428813>.
 - [19] Virginia Braun and Victoria Clarke. Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101, 2006. ISSN 1478-0887. doi: 10.1191/1478088706qp063oa. URL <https://doi.org/10.1191/1478088706qp063oa>.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [supplementary.pdf](#)