

Supplementary Information

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

A Survey Instruments

The instruments were administered as voluntary program evaluation with results reported only in aggregate; response options are abbreviated here.

A.1 Post-workshop evaluation

Five-point items unless noted: (1) overall satisfaction with the workshop; (2) relevance of the content to professional development; (3) clarity of the information presented; (4) effectiveness of the presenters; (5) extent to which the presenters created an interactive learning environment; (6) value of the hands-on exercises and breakout sessions; (7) topics for which the respondent had prior knowledge (select all); (8) topics found most engaging or valuable (select all); (9) topics desired in future workshops (open-ended); (10) likelihood to recommend the workshop to a colleague; (11) additional feedback (open-ended). Items 1–6 and 10 populate Table 5 of the main manuscript; item 7 records prior topic knowledge; items 8, 9, and 11 feed the qualitative analysis (Table 6 of the main manuscript).

A.2 Retrospective self-efficacy block (2026)

On the separate 2026 awareness form (distinct from the general post-workshop evaluation of Table 5 of the main manuscript), participants rated their confidence “before the workshop” and “now,” from 1 (not at all) to 5 (very), on ten threat topics (steganography and hidden-data detection; adversarial images; SQL injection; evil-twin Wi-Fi; wireless, NFC, and sub-GHz; USB and peripheral; AI/LLM phishing; deepfakes and synthetic identity; quantum and post-quantum; and account security) plus overall awareness of AI-era scams and a before/now item on having a practical sense of how attacks happen. The form also carried agreement items (1–5) on what the workshop taught, which threats could realistically affect the respondent, the pace of change, and where to keep learning, and on the hands-on activities, the speakers, safe and responsible delivery, and the five-day format.

A.3 Pre-workshop needs assessment

Prospective participants rated, from 1 (low) to 5 (very high), how strongly each of a broad set of cybersecurity topics should be emphasized, grouped into current threats, emerging and future threats, standards and frameworks, infrastructure and tools, and risk and compliance.

A.4 Expert-evaluator relevance panel

For each of the same ten threat topics, the expert evaluators rated (a) importance to the workforce and (b) current workforce preparedness, from 1 (very low) to 5 (very high), and answered open-ended items on missing threats, the defense-first approach, the workforce skills gap, and advice for adopters.

B Post-Workshop Response Distributions

Table S1 gives the full response distributions for the seven five-point post-workshop items summarized as favorable percentages in Table 5 of the main manuscript, reported per year over the same $N = 22$ valid responses per year (1 = least favorable, 5 = most favorable).

Table S1: Per-item response distributions for the seven post-workshop items, per year ($N = 22$ per year).

Item (5-point)	Year	1	2	3	4	5	N
Overall satisfaction	2025	1	0	0	3	18	22
	2026	0	0	0	7	15	22
Content relevance	2025	0	0	3	8	11	22
	2026	0	0	3	6	13	22
Clarity of material	2025	0	0	2	13	7	22
	2026	0	0	1	8	13	22
Presenter effectiveness	2025	0	0	3	9	10	22
	2026	0	1	1	7	13	22
Interactive learning	2025	0	0	1	5	16	22
	2026	0	0	1	7	14	22
Hands-on value	2025	0	1	2	4	15	22
	2026	0	0	2	8	12	22
Would recommend	2025	0	0	2	3	17	22
	2026	0	0	1	6	15	22

C Curriculum Modules

The five-day program ran 8:30 a.m.–4:30 p.m. each day, opening with check-in and breakfast and closing with interactive games and guided reflection, and was offered as two interchangeable weekly sessions each year. Each day also included an invited practitioner talk over lunch.

Each module below is described by its *aim* (what the learner should come to understand), its *format* expressed in the five-step teaching cycle of Section 3 of the main manuscript (Figure 3 there), and its intended *defensive outcome* (the takeaway of Table 1 of the main

manuscript). Three formats recur, reflecting the safety rules of Section 5 of the main manuscript: *full cycle* modules run all five steps including a guided hands-on lab; *demonstration-bounded* modules run steps 1, 2, 4, and 5, replacing the participant lab with an instructor-led demonstration, either because the technique fails the self-containment rule or because the awareness goal is reached by observation; and *briefing* modules run steps 1, 4, and 5 for topics whose awareness goal needs no practical component. The modules are grouped by day and map to the risk clusters of principle P1. Further information about the workshop, including the full agenda and program materials, is available on the program website (<https://inquirelab.ai/cybersecurity/>).

Day 1: Foundations and media trust

Introduction and emerging trends (briefing). *Aim:* orient participants to how the threat landscape is shifting and why non-specialists are now in scope. *Format:* opening briefing that frames the week, introduces the six risk clusters, and sets the defense-first stance; closing discussion connects the landscape to participants' own sectors. *Outcome:* a shared vocabulary of risk categories that the rest of the week's modules attach to.

Steganography and hidden-data detection (full cycle). *Aim:* understand how everyday media becomes an attack surface, with hidden payloads inside images, audio, and video that slip past file scanners. *Format:* concept and attacker's logic (why hiding data in media evades signature-based scanning), instructor demonstration of a payload concealed in an ordinary file, then a guided hands-on detection exercise on program devices; defensive framing and reflection close the session. *Outcome:* verify the authenticity of digital media and treat unexpected media files as a potential carrier, not just a document.

Password management and multi-factor authentication (full cycle). *Aim:* understand how credential theft succeeds and how password managers and multi-factor authentication break the attack chain. *Format:* concept and attacker's logic (credential reuse and phishing economics), demonstration of why weak and reused passwords fail, then guided hands-on setup of a password manager and multi-factor authentication. *Outcome:* hardened personal and account security habits that participants can apply the same day.

Day 2: Web and AI-application attacks

CyberArena: SQL injection, exploit and defend (full cycle). *Aim:* understand how unvalidated input becomes code execution, and why layered defenses matter. *Format:* concept and attacker's logic (why a login form that concatenates input is a program the attacker can rewrite), safe demonstration on a deliberately vulnerable login on a program-owned system, then a guided hands-on exercise in which participants run the same injection attempts against a machine-learning-defended version and see them fail; defensive framing contrasts the two systems. *Outcome:* recognize injection-class risk in web applications and the case for input validation plus defense-in-depth.

Retrieval-augmented-generation (RAG) attacks (demonstration-bounded). *Aim:* understand how AI assistants that pull from external data can be manipulated through poisoned or injected content. *Format:* concept and attacker's logic (the assistant trusts retrieved text as if the user wrote it), guided walkthrough of a manipulation scenario, then

the guardrails that limit the damage. *Outcome*: treat AI-assistant output as an untrusted input to decisions, and know the guardrails to ask for.

Cybersecurity in biological implants (briefing). *Aim*: recognize the security and privacy risks in networked medical and implantable devices. *Format*: briefing on the attack surface of safety-critical connected hardware, closing with defensive framing and discussion. *Outcome*: a defensive mindset toward safety-critical devices, in procurement and in use.

Vulnerability triage and patch strategy (full cycle). *Aim*: understand how to prioritize and remediate vulnerabilities under real-world constraints. *Format*: concept (how vulnerability disclosure floods defenders with alerts), then a guided hands-on triage-and-patch exercise turning a set of alerts into an ordered response. *Outcome*: prioritize patching by risk rather than by arrival order.

Day 3: AI-era and physical-layer threats

AI-powered attacks and defenses (full cycle; the worked example of Section 3.1 of the main manuscript). *Aim*: understand that machine-learning models can be fooled by inputs that look unchanged to a person, and why that matters wherever classifiers gate decisions. *Format*: the complete cycle traced in Section 3.1 of the main manuscript: concept, attacker’s logic (the fast gradient sign method), instructor demonstration on MobileNet, then a scaffolded notebook lab in which participants adjust the size of the change and watch the label and confidence change. *Outcome*: AI components need their own security testing; defenses include adversarial training, input validation, and a human in the loop for high-stakes decisions.

Quantum computing and post-quantum cryptography (full cycle). *Aim*: understand the “harvest now, decrypt later” risk and what the post-quantum transition asks of ordinary organizations. *Format*: concept and attacker’s logic (record encrypted traffic today, decrypt when quantum hardware allows), then a hands-on quantum key-exchange game that makes the cryptographic idea tangible, closing with the post-quantum cryptography already shipping in mainstream messaging and VPN tools. *Outcome*: recognize long-lived data as the asset at risk and cryptographic migration as a planning problem that starts now.

Digital Pickpockets: NFC and USB attacks (demonstration-bounded). *Aim*: understand how proximity interfaces (contactless cards, USB ports) become attack surfaces. *Format*: live instructor station that clones an access card with a hand-held radio tool and reveals the signals devices broadcast, alongside a keystroke-injection USB device; defensive framing covers the endpoint controls that stop it. Participants observe; no payload authoring is taught (Section 5 of the main manuscript). *Outcome*: treat unknown USB devices and unattended cards as live risks, and know the endpoint and physical controls that mitigate them.

Cybersecurity in smart cities and blockchain (briefing). *Aim*: recognize threats to connected urban infrastructure and distributed-ledger systems. *Format*: briefing on where these systems’ trust assumptions break down, with defensive framing and discussion. *Outcome*: evaluate connected-infrastructure and ledger claims by their trust assumptions rather than their novelty.

Day 4: Network defense and cryptography

DDoS mitigation, intrusion detection, and incident response (demonstration-bounded).

Aim: understand how denial-of-service attacks overwhelm a service, how intrusion-detection signals surface an attack already in progress, and the initial incident-response steps that follow. *Format:* concept and attacker's logic (exhausting a resource is cheaper than breaking a control), instructor demonstration of detection signals, defensive framing on mitigation, escalation, and incident response. *Outcome:* recognize the signals of an attack in progress, the reporting path they should trigger, and the initial response actions that follow detection.

The Invisible World: radio-frequency and infrared (demonstration-bounded).

Aim: make tangible the radio-frequency and infrared signals that everyday devices broadcast. *Format:* live demonstration revealing the invisible side of everyday wireless, with defensive framing on what those emissions expose. *Outcome:* awareness that devices communicate constantly and that the wireless boundary of an organization extends beyond its walls.

Cryptography and post-quantum migration (briefing). *Aim:* build the cryptography foundations needed to reason about encryption decisions, and understand what migrating everyday systems to post-quantum standards involves. *Format:* foundations briefing building on Day 3's key-exchange game, closing with migration planning and discussion. *Outcome:* reason about encryption and cryptographic readiness in one's own systems.

Understanding current cybersecurity threats (briefing). *Aim:* survey the threats most active now, from ransomware to AI-scaled phishing, and how they reach non-specialists. *Format:* briefing connecting the week's risk clusters to the current threat picture, with defensive framing and discussion. *Outcome:* place each emerging headline threat into a familiar risk category with a known defensive response.

Day 5: Tools, identity, and social engineering

Network-security-tools lab (full cycle). *Aim:* understand what defenders can see, and how the tools that map, monitor, and protect a network work. *Format:* guided hands-on lab with defensive network tools in the isolated program environment, with defensive framing on where each tool fits in a monitoring posture. *Outcome:* apply defensive tooling, and know what visibility to ask of one's own organization.

Synthetic identity and AI-driven social engineering (demonstration-bounded).

Aim: understand how attackers use AI to scale fake identities and convincing scams. *Format:* concept and attacker's logic (generation makes impersonation cheap; the target's verification habit is the control), demonstration of how synthetic content presents, then the verification habits that defeat it. *Outcome:* verification habits (out-of-band confirmation, challenge questions) that hold even when the deception is high quality.

Wi-Fi (evil-twin) network defense (demonstration-bounded). *Aim:* understand how rogue access points impersonate trusted Wi-Fi to intercept traffic. *Format:* concept and attacker's logic (a network name is a claim, not an identity), bounded instructor demonstration of how the deception presents to a user, then the habits that cut the risk fastest; no interception of real traffic and no payload authoring (Section 5 of the main manuscript). *Outcome:* habits for joining networks defensively and recognizing an impersonated access point.

Identity-and-access management with a hardware root of trust (briefing). *Aim:* understand how identity and access are managed at scale, and how a hardware root of trust (a trusted platform module) anchors device and identity security. *Format:* briefing connecting Day 1's account-security habits to organizational identity architecture, with defensive framing and discussion. *Outcome:* manage identity and trust; know why hardware-anchored identity resists the credential attacks of Days 1 and 5.