Designing Cybersecurity Education as a Socio-Technical Capability: The CEDDIE Framework

Pilleriin Sara Lillemets[1]

[1]Department of Applied Mathematics and Computer Science (DTU Compute), Technical University of Denmark, Richard Petersens Plads, Building 322, 2800 Kgs. Lyngby, Denmark

Abstract

This paper presents the doctoral research developing CEDDIE, a socio-technical framework for designing, implementing, and continuously adapting cybersecurity education and training as an operational capability. The framework extends the established ADDIE instructional design lifecycle by replacing the Analysis phase with an Adapt phase in subsequent cycles, converting training from a repeating project into a continuous capability loop. The research uses theoretical perspectives from socio-technical systems theory, dynamic capabilities theory, and organisational learning theory to generate specific design requirements that the framework's architecture addresses. So far, three Design-Based Research cycles have demonstrated the feasibility and utility of the Analysis phase. The central proposition that the full cycle functions as a continuous adaptation mechanism remains to be tested through planned full-cycle evaluation across multiple institutional contexts.

Keywords

cybersecurity education, training needs analysis, socio-technical systems, ADDIE, competency frameworks, digital twins, cyber ranges

1. Introduction

As organisations adopt new digital technologies, they simultaneously acquire new vulnerabilities, new coordination requirements, and new failure modes. This process of digital transformation has consequences for how cybersecurity education and training should be designed. When an organisation's operating model is stable, periodic training can be sufficient. When it changes continuously, however, periodic training becomes structurally inadequate: the gap between required and trained capability widens between deliveries, and each technology adoption or process change introduces misalignments that the training system cannot detect or correct.

This doctoral research addresses that mismatch. The central claim is that, under conditions of continuous digital transformation, cybersecurity education and training must be designed and governed as an operational capability that is continuously adapted rather than delivered as a sequence of disconnected events. This requires mechanisms that translate changing operational demands into role-relevant capability requirements, connect those requirements to defensible training design decisions, and feed evidence from delivery back into subsequent iterations.

To that end, the research develops a socio-technical framework for cybersecurity education and training called CEDDIE. The name combines "Cyber Education" with ADDIE, the established instructional design method of Analysis, Design, Development, Implementation, and Evaluation. CEDDIE retains this five-phase structure but introduces a structural modification: after the initial cycle, which begins with Analysis, the first phase becomes Adapt in consecutive cycles. The Adapt phase drives targeted revisions to training design based on evaluation evidence and the operational environment. This modification turns the training lifecycle from a repeating project into a continuous capability loop, where each delivery generates evidence that feeds the next adaptation decision.

The framework is grounded in theoretical perspectives from socio-technical systems theory, dynamic capabilities theory, and organisational learning theory, generating combined design requirements for a

Baltic DB&IS 2026 Conference Forum and Doctoral Consortium, 28 June - 1 July 2026, Tartu, Estonia

pilli@dtu.dk (P. S. Lillemets)

0009-0001-7003-5175 (P. S. Lillemets)

training system operating under continuous change. These are elaborated in Section 2. Throughout the doctoral project, this framework will be constructed and empirically tested in a Design Science Research programme comprising systematic reviews, Design-Based Research cycles conducted with training providers, and an eventual full-cycle evaluation in different institutional contexts.

The doctoral project is guided by one overarching research question:

> *What design principles must a socio-technical framework satisfy to sustain cybersecurity education and training as a continuous, evidence-driven operational capability in fast-changing environments?*

This is decomposed into three sub-questions:

RQ1. How can cybersecurity roles, tasks, and organisational needs be integrated with learning objectives, proficiency targets, and assessment criteria?

RQ2. How can cyber ranges and digital twins support repeatable and adaptive training interventions?

RQ3. How should training interventions be evaluated and refined to support alignment with operational needs and organisational learning?

The remainder of this paper is structured as follows. Section 2 presents the theoretical foundations and the design requirements they generate. Section 3 reviews the state of the art across instructional design, competency frameworks, and self-assessment. Section 4 describes the CEDDIE framework. Section 5 outlines the methodology and empirical programme. Section 6 summarises contributions and current status. Section 7 discusses open issues, and Section 8 concludes.

2. Theoretical foundations

The framework draws on three theoretical perspectives. Each is presented here as a source of specific design requirements: claims about what a cybersecurity training system must do if it is to function as a continuous capability under conditions of organisational change.

2.1. Socio-technical systems: training a changing organisation

Socio-technical systems theory holds that organisational performance is an emergent property of the interaction between social and technical subsystems, and that optimising either subsystem independently degrades the performance of the whole [1]. Applied to cybersecurity, this means that secure practice cannot be achieved by training individuals in technical procedures alone, because operational performance depends on coordination structures, shared situation awareness, and distributed decision-making that span roles and organisational boundaries.

For cybersecurity training, this generates a specific requirement: the training needs analysis must capture coordination and decision-making demands as well as individual technical competencies. A training system that analyses needs only at the level of individual technical skill will systematically miss the coordination demands that determine operational effectiveness.

The CEDDIE framework responds to this requirement by integrating competency frameworks with proficiency calibration. The NICE Workforce Framework [2] provides the task and role vocabulary, but it is intentionally proficiency-agnostic, as selecting a task does not specify the cognitive demand expected. Bloom's Revised Taxonomy [3] provides the missing depth dimension. This integration is relevant because the same competency performed at different cognitive levels represents fundamentally different operational demands: executing an incident response procedure (Apply) and judging which containment strategy is appropriate under time pressure and incomplete information (Evaluate) are different tasks from a socio-technical perspective, even when they address the same domain content. Without this depth specification, a training needs analysis cannot distinguish between procedural gaps and decision-making gaps: a procedural gap requires guided practice, while a decision-making gap requires scenario conditions that force judgement under uncertainty.

2.2. Dynamic capabilities: sensing, seizing, and reconfiguring

Dynamic capabilities theory argues that in changing environments, sustained performance depends on the organisation's capacity to sense emerging demands, seize opportunities by translating them into actionable responses, and reconfigure routines accordingly [4]. In this view, an organisation's operational capabilities are the routinised functions it performs reliably as part of ongoing operations; dynamic capabilities are the higher-order capabilities that adapt them.

Applied to cybersecurity training, this predicts that generic training delivered on a fixed schedule will progressively drift from operational reality. The training function itself must operate as a dynamic capability. Sensing requires detecting changes in the operational environment, such as threat landscape shifts, technology adoption, dependency changes, and regulatory developments. Seizing requires translating those detected changes into concrete training design decisions within a timeframe that preserves operational relevance. Reconfiguring requires updating delivery routines, materials, and assessment criteria without full redesign. The combination is what this research defines as continuous adaptation.

CEDDIE's Adapt phase is foreseen to be the reconfiguration mechanism, suggesting that evaluation evidence from the previous cycle and signals from the operational environment drive targeted revisions. The Analysis phase constitutes a seizing mechanism, translating competency requirements into gap profiles and scenario configurations. However, sensing is the least developed component. The Adapt phase is designed to draw on both backward-looking evidence (from evaluation) and forward-looking signals (from the operational environment). Currently, the latter input is specified at the framework level without an empirically validated mechanism for producing it. Digital twins are investigated as a candidate for this sensing function, as a continuously updated operational representation that can provide the upstream signal the Adapt phase requires.

2.3. Organisational learning: single-loop and double-loop

Organisational learning theory distinguishes between individual learning and the embedding of that learning in shared routines, mental models, and governance structures that persist beyond the individuals who generated them [5, 6]. It further distinguishes single-loop learning - detecting and correcting errors within existing frames of reference - from double-loop learning, which involves questioning and revising the frames themselves.

In cybersecurity contexts, most training systems operate in single-loop mode: the same training or exercise is delivered, minor content updates are made in response to participant feedback or new threat intelligence, and the cycle repeats. This is adequate when the training design assumptions are sound, but it cannot detect or correct systematic misalignments. For instance, training that consistently targets the wrong cognitive level, addresses the wrong competency mix, or assumes coordination structures that do not exist in participants' actual work environment. Double-loop learning requires that the training system can interrogate its own design assumptions, not only its content.

This generates two requirements. First, evaluation must produce evidence about design assumptions, such as gap profile accuracy, scaffolding effectiveness and competency area appropriateness. Second, design knowledge must be captured in ways that are institutionally persistent, meaning that an organisation can maintain its training capability without dependence on the original designer. CEDDIE addresses the first requirement through its iterative structure. It addresses the second by standardising the process through which design knowledge is captured, providing a common set of procedures and data structures that enable different designers to work from the same evidence base and follow the same design logic.

3. State of the art

This doctoral research draws on three bodies of literature: instructional design for cybersecurity education, competency and proficiency frameworks, and self-assessment in professional training. These

strands are typically treated separately. Instructional design research addresses how training should be structured and delivered. Competency frameworks specify what should be trained. Self-assessment research examines how learner readiness can be approximated under practical constraints. The research gap lies at their intersection: the translation from workforce-relevant competency requirements through cohort-specific readiness diagnostics to defensible, traceable training design decisions.

3.1. Instructional design for cybersecurity education and training

Developed in the 1970s as the U.S. military's Interservice Procedures for Instructional Systems Development (IPISD) [7], the ADDIE model has become widely adopted, extending its use from military training to broader applications, including cyber education and training. However, instructional design scholarship cautions that the model is better understood as a flexible, iterative family of practices than as a rigid linear sequence [7]. Recent refinements emphasise rapid formative feedback loops and the application of personalised learning theory to game-based and tabletop cybersecurity exercises [8, 9]. These establish a defensible baseline for applying ADDIE to cybersecurity training while highlighting the distance between lifecycle guidance and delivery practice.

A recurring theme is that instructional design decisions are tightly coupled to the affordances and constraints of delivery infrastructure. Research on cyber range and testbed design emphasises that realism and repeatability depend on deliberate architectural choices such as isolation, instrumentation, controllability, and that role-focused scenario engineering can improve alignment with workforce needs but introduces operational overhead that limits reuse across cohorts [10, 11]. These platform-driven constraints make the Analysis phase especially consequential: when cohort readiness and performance targets are not articulated early, adaptation is deferred to delivery, where it is harder to justify and more costly to enact.

3.2. Competency frameworks and proficiency calibration

Competency frameworks provide the shared vocabulary for connecting workforce needs to educational design. The NICE Workforce Framework for Cybersecurity specifies work through Task, Knowledge, and Skill statements organised into Work Roles and Competency Areas [2, 12]. In the European context, the ENISA European Cybersecurity Skills Framework plays a comparable role [13, 14]. Both provide a systematic way of describing what cybersecurity professionals do, but neither provides a built-in proficiency metric. NIST IR 8355 explicitly recommends omitting proficiency qualifiers from competency descriptions so that the same competencies can be applied across proficiency levels [12]. This is a deliberate and reasonable design choice, but it creates a practical gap for instructional design: selecting a task does not by itself specify the intended cognitive demand, the expected depth of performance, or the assessment standard.

Bloom's Revised Taxonomy offers a complementary lens. Krathwohl's [3] revision distinguishes between a cognitive-process dimension and a knowledge dimension and retains a broadly hierarchical progression from remembering through creating. In security education, Bloom's has been used to calibrate training to different audiences [15], to structure curricula around increasingly demanding outcomes [16], and to inform competency assessment rubrics [17]. The value of Bloom's in this context is that it helps distinguish tasks requiring routine application from those requiring analytic interpretation or evaluative judgement under pressure. This is a distinction that matters considerably in cyber training where tasks may share domain content but differ in cognitive demand.

3.3. Self-assessment and readiness instruments

In time-constrained professional training, readiness is often approximated through self-assessment. The literature consistently shows that self-rated competence correlates weakly with externally observed performance, particularly when learners make broad, unanchored judgements about their own ability [18, 19]. However, self-assessment becomes more informative when anchored to concrete tasks and explicit standards [20]. This suggests that pre-exercise questionnaires are most defensible when treated as

self-reported readiness indicators rather than as direct competence measures, which in turn determines how the resulting data should be interpreted and used. Recency of practice is relevant as a separate dimension. Skills decay with non-use, but decay rates vary by task type, complexity, and opportunities for intermittent practice [21]. This means recency is better treated as a contextual indicator.

3.4. Research gap

Taken together, the literature points to a gap at the interface between these three strands. Competency frameworks provide the vocabulary for specifying what should be trained but not the depth at which performance is expected. Bloom's Revised Taxonomy provides the depth lens but is not integrated into standard workforce frameworks. Self-assessment research identifies the conditions under which pre-exercise instruments are defensible but does not connect these assessments to scenario configuration decisions. The result is that the translation from organisational need through cohort readiness to training design remains largely ad hoc and dependent on individual instructor judgement rather than on a structured, traceable process. The CEDDIE Framework is proposed as a response to this gap.

4. The CEDDIE framework

CEDDIE is a socio-technical framework for designing, implementing, and continuously adapting cybersecurity education and training as an operational capability. It inherits the five-phase structure of ADDIE - Analysis, Design, Development, Implementation, and Evaluation - but modifies it in one structurally consequential way: after the initial cycle, which begins with Analysis, subsequent cycles begin with Adapt. In the Adapt phase, evaluation evidence from the previous iteration and signals from the operational environment drive targeted revisions to training design, replacing the full analytical effort of the first cycle with an evidence-informed recalibration. This is the mechanism that converts the lifecycle into a continuous capability loop.

The framework integrates two complementary alignment mechanisms. The NICE Workforce Framework for Cybersecurity provides a role- and task-based vocabulary for specifying the content of training: what competencies, tasks, and knowledge areas a given training audience needs to develop [2]. Bloom's Revised Taxonomy provides a depth lens for specifying to what cognitive level performance is expected [3]. Neither mechanism is sufficient alone. NICE without Bloom's specifies tasks but not the standard of performance; Bloom's without NICE provides a cognitive demand hierarchy but no connection to workforce-relevant role descriptions. Together, they enable training design to be grounded in both workforce relevance and defensible performance expectations.

The five phases operate as follows at the framework level. Analysis characterises the training need: it identifies the target roles and competency areas, specifies proficiency targets using Bloom's levels, elicits cohort readiness through a structured self-assessment instrument, and produces a gap profile that quantifies the distance between target and current proficiency for each competency area. Design translates the gap profile into training design decisions: scenario selection, learning objective specification, scaffolding tier assignment, and the configuration of instructional support mapped to specific scenario steps. Development produces the training materials, scenario configurations, and instructor guidance required for delivery. Implementation is the delivery itself: the conducted exercise or training event. Evaluation captures evidence from delivery: trainee performance, instructor observations, scaffolding effectiveness, and design assumption validity. In subsequent cycles, Adapt replaces Analysis as the opening phase. Adapt draws on two inputs. The first is evaluation evidence from the previous cycle - what worked, where gaps persisted, which scaffolding decisions were effective, and which design assumptions proved correct or incorrect. The second is a current representation of the operational environment, such as changes in the organisation's technology stack, threat exposure, dependency structure, or regulatory obligations that affect what training should address. The first input enables single-loop learning: correcting content and configuration within the existing design frame. The second input, combined with evidence about design assumption validity from Evaluation, enables double-loop learning: revising the frame itself when the operational context has shifted sufficiently to warrant it.

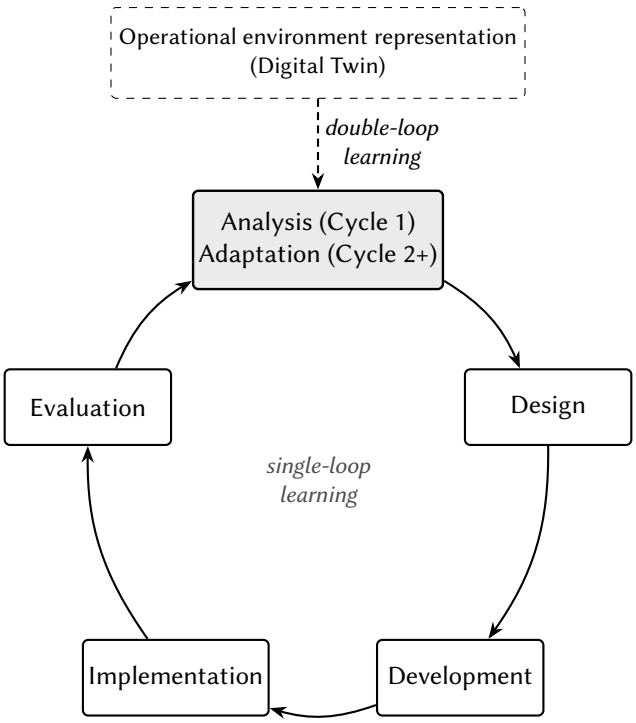

Figure 1: The CEDDIE lifecycle. Cycle 1 begins with Analysis; from Cycle 2 onwards, Adapt replaces Analysis as the opening phase.

The integration of Bloom's-calibrated proficiency targets with NICE competency specifications addresses the socio-technical requirement: training needs are characterised across cognitive levels, making both procedural and coordination demands visible within the same analysis. The Adapt phase functions as the reconfiguration mechanism required by dynamic capabilities theory, while the TNA instrument and Training Design Matrix (described in Section 6) together constitute the seizing mechanism that translates sensed demands into concrete design decisions. The evaluation-to-adaptation feedback loop, structured to capture evidence about design assumptions as well as trainee performance, provides the basis for the double-loop learning that organisational learning theory requires.

One theoretical requirement is not yet fully addressed. Dynamic capabilities theory requires a forward-looking sensing mechanism that detects operationally relevant changes before they manifest as training failures. The Adapt phase is designed to incorporate such signals through its second input (operational environment representation), but no validated mechanism for producing this input currently exists. Cybersecurity digital twins are under investigation as a candidate, and this is a central question for the planned empirical work described in Section 5.3.

5. Methodology

5.1. Overarching approach

The doctoral research follows a Design Science Research (DSR) approach, in which the primary knowledge contribution is a purposefully constructed artefact - the CEDDIE framework and its associated instruments - together with the design knowledge generated through building, testing, and refining that artefact in practice [22]. Within this overarching approach, the research employs multiple methods matched to different research questions and settings: systematic literature reviews for knowledge base construction, Design-Based Research (DBR) for iterative empirical testing in authentic training environments [23], qualitative methods (semi-structured interviews and non-interventionist observation) for exploratory and explanatory work, and descriptive quantitative analysis for cohort diagnostics.

Table 1
Design-Based Research cycles.

Cycle	Cohort type	n	Scale	Key outcome
1	Municipal IT/security	15	Dreyfus 1–5	Baseline TNA; scale mismatch and item misfit
2	National-level agency	8	Dreyfus 1–5	Confirmed mismatch; combined-score problem
3	Public-sector cohort	8	Bloom's 0–6	Resolved mismatch; separated recency; NICE-traced items

5.2. Empirical cycles

The core empirical work to date comprises three DBR cycles conducted in collaboration with a cyber range training provider. All three cycles used the same ransomware investigation scenario with different public-sector cohorts. The scenario required participants to investigate a ransomware intrusion using SIEM log analysis, digital forensics, ATT&CK-based threat analysis, and coordinated incident response.

The training needs analysis instrument was refined across cycles in response to three identified problems. First, a construct mismatch between Dreyfus-derived self-assessment scales and Bloom's-informed proficiency targets, resolved in Cycle 3 by adopting a unified Bloom's-aligned scale. Second, the conflation of proficiency and recency in a composite score lacking empirical justification, resolved by separating recency as a distinct contextual indicator. Third, item wording tied to specific tools rather than to task-level competencies, resolved by deriving each item from a scenario learning objective with a documented crosswalk to NICE Task, Knowledge, and Skill statements.

Data were gathered from pre-exercise TNA questionnaires, non-interventionist observation during the exercise delivery, and semi-structured interviews with trainees and instructors across the first two cycles. Quantitative analysis was descriptive: item-level gaps (target Bloom's level minus self-assessed proficiency) were aggregated by competency area to produce cohort-level readiness profiles. Qualitative analysis used directed coding in Atlas.ti. Because the scales are ordinal and the samples small, the analysis is interpretive rather than inferential.

5.3. Planned empirical work

The empirical work to date has tested the Analysis phase. The remaining phases are specified at the framework level but lack equivalent empirical grounding. Next phases of the research are designed to address this gap across two settings.

The planned empirical work is structured around a five-month external research stay in 2026. A one-month visit to a mature operational cyber range in a defence training context will examine how training needs identification and scenario currency are handled in practice, establishing a baseline against which CEDDIE's structured approach can be compared and testing transferability beyond the initial training-provider setting. From September to December 2026, a four-month stay at NICS Lab, University of Málaga, will investigate cybersecurity digital twins as a candidate sensing mechanism for the Adapt phase, addressing the forward-looking sensing requirement identified in Sections 2 and 4.

Together, these settings should extend the evidence base beyond the single-provider, single-scenario, single-national-context limitations of the initial study. They also address specific theoretical questions: whether the full cycle enables the continuous reconfiguration that dynamic capabilities theory requires, and whether the design knowledge captured through iteration can function as institutional memory in organisations other than the one in which it was developed.

6. Contributions and current status

The central contribution of this doctoral research is a socio-technical framework for designing, implementing, and continuously adapting cybersecurity education and training as an operational capability, supported by empirical research that progressively constructs, tests, and refines its constituent elements.

First, it reframes cybersecurity education as a governed, continuous function rather than a sequence of discrete events, grounding this in socio-technical systems theory, dynamic capabilities theory, and organisational learning theory. Second, it integrates two alignment mechanisms that are typically treated separately in the literature: NICE for workforce-relevant content specification and Bloom's Revised Taxonomy for cognitive demand calibration. Third, it introduces the Analysis-to-Adapt structural modification that converts the ADDIE method from a repeating project into a continuous capability loop with explicit provisions for both single-loop and double-loop learning.

At the artefact level, the empirical research has so far produced three outputs. The TNA instrument implements the Analysis phase as a structured self-assessment aligned to NICE competency specifications and calibrated using Bloom's Revised Taxonomy, generating cohort-level readiness profiles. The Training Design Matrix connects those profiles to scenario scaffolding decisions, making the transition from analysis to design traceable and repeatable. The instrument evolution across three DBR cycles documents the construct-level problems identified through empirical use and the revisions made in response.

The doctoral project is planned to run from August 2024 to July 2027. The three Design-Based Research cycles reported in this paper were conducted between late 2025 and spring 2026. A journal article reporting the detailed findings is in preparation. In addition, the doctoral research has to date produced a co-authored systematic review of cyber range taxonomies [24] and a co-authored comparative analysis of open-source cyber ranges [25]. A workshop paper extending the TNA approach into during-exercise observation and post-exercise assessment is under review at the IEEE CSR 2026 CRIRM Workshop.

Considering the state of the research project, the scope of current claims is deliberately bounded. Evidence to date supports the feasibility, interpretability, and design utility of the Analysis phase under authentic delivery constraints. It does not yet demonstrate effects on learning outcomes, transfer to operational settings, or sustained organisational capability improvement. The broader claim that the CEDDIE Framework can be used to create a continuous operational capability through the Adapt mechanism remains a theoretically grounded proposition. Testing it is the purpose of the planned empirical work.

7. Open issues

7.1. From analysis to full cycle

The most consequential gap is that only the Analysis phase has been empirically tested. The remaining phases are specified at the framework level but lack equivalent empirical grounding. This matters because the distinctive claim of CEDDIE is not that the Analysis phase works, but that the full cycle functions as a continuous adaptation mechanism. Until the Adapt loop has been observed and documented, the central proposition remains theoretical. The external research stay period is designed to address this, but the practical and methodological challenges of testing a full lifecycle in authentic settings should not be underestimated.

7.2. Sensing and forward-looking adaptation

As identified in Section 2, dynamic capabilities theory requires that the Adapt phase draws on forward-looking signals about the operational environment. The Adapt phase is designed to incorporate such signals through its second input (operational environment representation), but no validated mechanism for producing this input for training design currently exists. Without such a mechanism, the framework risks systematic lag, revising training on the basis of the last delivery rather than aligning it to the present operational state. Digital twins are under investigation as a candidate: a continuously updated operational representation that can provide an upstream signal for Adapt, distinct from cyber ranges, which serve as training delivery infrastructure. Whether such representations can realistically provide a

reliable signal given fidelity constraints, maintenance costs, and governance requirements is the central question for the planned work.

7.3. Measurement limitations

The TNA instrument's reliance on self-assessed proficiency is a limitation, as self-reports correlate inconsistently with observed performance. While currently used for scenario configuration, the framework's maturation requires triangulation with objective behavioural evidence to ensure robust evaluation without increasing instructor burden. Furthermore, the Bloom's-aligned 0–6 scale, while resolving previous construct mismatches, was not designed as an interval scale for self-assessment. The consistency of respondent interpretation and the reliability of resulting gap values for instructional decision-making remain critical questions for future methodological investigation.

8. Conclusion

This paper has presented the ongoing doctoral research developing CEDDIE, a socio-technical framework for designing, implementing, and continuously adapting cybersecurity education and training as an operational capability. The empirical research so far has demonstrated the feasibility and design utility of the Analysis phase through three Design-Based Research cycles. The central proposition that the full framework can function as a continuous adaptation mechanism remains to be tested through the planned full-cycle evaluation across multiple institutional contexts and delivery formats. The boundary between what has been demonstrated and what has been claimed is maintained deliberately: this research contributes a framework that is theoretically motivated, partially validated, and designed for the empirical scrutiny that the next phase of work will provide.

Acknowledgments

This doctoral research is conducted at DTU Compute, Technical University of Denmark, under the supervision of Prof. Nicola Dragoni. The project is part of a collaboration between DTU and the Danish Defence, aimed at strengthening Denmark's digital resilience through research on cybersecurity education and training. The author gratefully acknowledges the training provider and the participants and instructors who contributed their time and expertise to the empirical cycles reported in this paper.

Declaration on Generative AI

During the preparation of this work, the author used Claude Opus 4.7 (Anthropic) to: improve writing style, paraphrase and reword, content enhancement, formatting assistance, peer review simulation. After using these tools/services, the author reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] F. E. Emery, E. L. Trist, Socio-technical systems, in: C. W. Churchman, M. Verhulst (Eds.), Management Sciences: Models and Techniques, volume 2, Pergamon Press, 1960, pp. 83–97.

[2] R. Petersen, D. Santos, M. C. Smith, G. A. Witte, K. Wetzel, Workforce Framework for Cybersecurity (NICE Framework), Technical Report NIST SP 800-181 Rev. 1, National Institute of Standards and Technology, Gaithersburg, MD, 2020.

[3] D. R. Krathwohl, A revision of Bloom's taxonomy: An overview, Theory into Practice 41 (2002) 212–218.

[4] D. J. Teece, Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance, Strategic Management Journal 28 (2007) 1319–1350.

[5] C. Argyris, D. A. Schön, Organizational Learning: A Theory of Action Perspective, Addison-Wesley, Reading, MA, 1978.
[6] D. H. Kim, The link between individual and organizational learning, Sloan Management Review 35 (1993) 37–50.
[7] M. Molenda, In search of the elusive ADDIE model, Performance Improvement 42 (2003) 34–37.
[8] N. Chowdhury, S. Katsikas, V. Gkioulos, Modeling effective cybersecurity training frameworks: A delphi method-based study, Computers & Security 113 (2022) 102551.
[9] N. Chowdhury, V. Gkioulos, A personalized learning theory-based cyber-security training exercise, International Journal of Information Security 22 (2023) 1531–1546.
[10] M. Frank, M. Leitner, T. Pahi, Design considerations for cyber security testbeds: A case study on a cyber security testbed for education, in: 2017 IEEE 15th Intl Conf on Dependable, Autonomic and Secure Computing (DASC/PiCom/DataCom/CyberSciTech), Orlando, FL, 2017, pp. 38–46.
[11] S. Karagiannis, E. Magkos, E. Karavaras, A. Karnavas, M. N. Nikiforos, C. Ntantogian, Towards NICE-by-design cybersecurity learning environments: A cyber range for SOC teams, Journal of Network and Systems Management 32 (2024) 42.
[12] K. Wetzel, NICE Competency Areas, Technical Report NIST IR 8355, National Institute of Standards and Technology, Gaithersburg, MD, 2023.
[13] ENISA, ECSF — European Cybersecurity Skills Framework Role Profiles, Technical Report, European Union Agency for Cybersecurity, Athens, 2022.
[14] ENISA, ECSF — European Cybersecurity Skills Framework User Manual, Technical Report, European Union Agency for Cybersecurity, Athens, 2022.
[15] J. F. van Niekerk, R. von Solms, Using Bloom's taxonomy for information security education, in: R. C. Dodge, L. Futcher (Eds.), Information Assurance and Security Education and Training, volume 406 of *IFIP Advances in Information and Communication Technology*, Springer, 2013, pp. 280–287.
[16] M. A. Harris, K. P. Patten, Using Bloom's and Webb's taxonomies to integrate emerging cybersecurity topics into a computing curriculum, Journal of Information Systems Education 26 (2015) 219–234.
[17] N. K. Ramsoonder, S. Kinnoo, A. J. Griffin, C. Valli, N. F. Johnson, Optimizing cyber security education: Implementation of Bloom's taxonomy for future cyber security workforce, in: 2020 International Conference on Computational Science and Computational Intelligence (CSCI), IEEE, 2020, pp. 93–98.
[18] M. J. Gordon, A review of the validity and accuracy of self-assessments in health professions training, Academic Medicine 66 (1991) 762–769.
[19] D. A. Davis, P. E. Mazmanian, M. Fordis, R. Van Harrison, K. E. Thorpe, L. Perrier, Accuracy of physician self-assessment compared with observed measures of competence: A systematic review, JAMA 296 (2006) 1094–1102.
[20] S. C. Hawkins, A. Osborne, S. J. Schofield, D. J. Pournaras, J. F. Chester, Improving the accuracy of self-assessment of practical clinical skills using video feedback: The importance of including benchmarks, Medical Teacher 34 (2012) 279–284.
[21] C. E. Tatel, P. L. Ackerman, Procedural skill retention and decay: A meta-analytic review, Psychological Bulletin 151 (2025) 696–736.
[22] A. R. Hevner, S. T. March, J. Park, S. Ram, Design science in information systems research, MIS Quarterly 28 (2004) 75–105.
[23] F. Wang, M. J. Hannafin, Design-based research and technology-enhanced learning environments, Educational Technology Research and Development 53 (2005) 5–23.
[24] P. Lillemets, N. B. Jawad, J. Kashi, A. Sabah, N. Dragoni, A Systematic Review of Cyber Range Taxonomies: Trends, Gaps, and a Proposed Taxonomy, Future Internet 17 (2025) 259.
[25] E. J. Lundqvist, P. S. Lillemets, N. Dragoni, A Comparative Analysis of Open-Source Cyber Ranges for Cyber-security Education, Procedia Computer Science 272 (2025) 415–420.