

Enhancing Data Privacy and Ethics Instruction in Information Security Programs

Robert Riley¹ (rriley@students.vfairmontstate.edu), Mark Reeves¹ (mreeves@students.vfairmontstate.edu)

¹College of Science and Technology, Fairmont State University, Fairmont, WV, United States

Abstract

This paper reports on the design, deployment, and evaluation of generative AI-assisted learning scenarios for undergraduate instruction in data privacy and information security ethics. Drawing on constructivist learning theory, a collaborative workflow between faculty and a large language model produced a set of scenario-based modules that address privacy regulation, organizational data governance, incident-response ethics, and surveillance accountability. Each scenario was iteratively refined based on input from a student advisory panel consisting of two graduate research assistants and one undergraduate research assistant. A pilot evaluation with 21 enrolled students employed a mixed-methods protocol that combined five-point Likert-scale items with open-ended written responses. Quantitative results showed consistently favorable ratings: 95.2% of respondents agreed or strongly agreed that the scenarios captured their interest, while 90.5% reported enhanced understanding of core privacy and ethics concepts. Thematic analysis of qualitative data revealed that students valued the multi-perspective document formats, the contemporary relevance of the situations depicted, and the realism of embedded artifacts such as simulated regulatory correspondence and internal memoranda. Challenges included occasional vocabulary complexity exceeding the expected level for introductory students, addressed through supplemental glossaries. The findings support the viability of human–AI collaborative content development as a scalable method for producing contextualized, scenario-driven instructional materials in the information security domain. Implications for curriculum design, faculty adoption, and future controlled studies are discussed. The scenario-development workflow, survey instruments, and interpretive claims are still being refined. Subsequent revisions will strengthen the empirical framing and expand methodological detail.

Keywords — Scenario-based learning; Generative AI; Data privacy education; Information security ethics; Curriculum development

1 Introduction

As data-driven systems become routine in business, healthcare, government, and education, information security graduates need both technical preparation and practice with the ethical and legal dimensions of data handling [1, 2]. One widely cited workforce report estimated a global cybersecurity workforce gap of 3.4 million practitioners [2]. For university programs, the gap is not only a matter of technical coverage; graduates also need repeated practice with regulatory compliance, privacy stewardship, and ethical judgment [3, 4]. Although cybersecurity offerings have expanded, ethics and data privacy often receive less sustained attention than technical topics [5, 6].

Scenario-based and case-study pedagogies have long been recognized as effective strategies for cultivating applied reasoning in professional domains [7, 8]. Within information security education specifically, scenario exercises that immerse students in realistic decision-making contexts have shown promise in developing both technical and ethical competencies [9–11]. The practical difficulty is production time. A useful scenario has to be plausible, include artifacts that feel authentic, reflect current regulatory language, and fit the students' level. Instructors in the present study treated that design work as the bottleneck that generative AI might help reduce; the broader literature likewise notes the cost of developing realistic cybersecurity training exercises [12].

The emergence of large language models (LLMs) capable of generating coherent, contextually appropriate text has introduced new possibilities for alleviating this bottleneck [13, 14]. Rather than replacing the instructor, these models can serve as collaborative partners in a structured content-creation workflow, producing draft scenarios that human experts subsequently refine and validate [15]. This human–AI collaborative paradigm aligns with constructivist principles that emphasize the active construction of knowledge through authentic, situated experiences [16, 17].

This paper reports a pilot study in which a faculty-led team used a generative AI model to draft scenario-based learning materials for an undergraduate course on data privacy and information security ethics. The course, housed within an information security degree program and aligned with national workforce framework categories [4], addresses six topical modules spanning classical ethical theory as applied to digital contexts, privacy regulation, organizational governance, incident response, surveillance and civil liberties, and the evolving landscape of AI governance. Each module received a suite of scenarios designed to provoke critical analysis of multi-stakeholder dilemmas. In this paper, we present an iterative workflow for human–AI scenario development, documenting design decisions, prompt engineering strategies, and revision cycles. Also, we report findings from a pilot deployment ($n = 21$), including descriptive survey results and a thematic analysis of open-ended responses.

Indeed, the reported data come from an initial pilot implementation, and several components, including comparison conditions, direct learning assessments, and expanded participant samples, remain under development.

2 Related Work

Cybersecurity education has repeatedly been framed as more than technical instruction; ethical, legal, and policy dimensions are part of professional preparation [3, 18]. Tavani [19] and Quinn [20] have articulated comprehensive ethical frameworks for computing professionals, yet the translation of these frameworks into classroom practice remains uneven. Multiple studies have documented the relative scarcity of dedicated ethics coursework in cybersecurity degree programs, noting that ethical content is often relegated to brief modules in technically oriented courses [5, 6].

Data privacy instruction poses a further challenge because regulatory environments change and because students must learn how to reason under uncertainty. The enactment of the General Data Protection Regulation (GDPR) in Europe, the California Consumer Privacy Act (CCPA), and analogous statutes in other jurisdictions has created a moving target for curriculum designers [21]. Students need exposure not just to the letter of current law but also to the reasoning processes by which professionals navigate ambiguous situations in which legal mandates conflict with organizational incentives or user expectations.

Constructivist learning theory posits that learners build understanding through active engagement with authentic problems embedded in meaningful contexts [8, 16]. Scenario-based learning operationalizes this principle by presenting learners with narrative situations that require analysis, judgment, and decision-making [7]. Within cybersecurity education, scenario-based approaches have been deployed for penetration testing exercises [9], incident response drills [10], and ethics instruction [11].

Scenario-based learning depends on several design factors: fidelity to real-world conditions, integration of multiple stakeholder perspectives, graduated complexity, and alignment with specific learning objectives [22]. Poorly constructed scenarios, which are often overly simplified, culturally narrow, or disconnected from course goals, can undermine student engagement rather than enhance it.

Game-based and interactive learning methods have also gained traction as strategies for boosting motivation and participation in cybersecurity courses [23]. These approaches share with scenario-based learning an emphasis on active participation, though they often prioritize procedural skill development over the ethical reasoning that is central to this study.

The public release of capable LLMs has prompted debate about their role in education [13, 14]. Much of this discussion has centered on student use, including concerns about plagiarism, assessment integrity, and the development of critical AI literacy [15]. Less attention has been paid to the instructor's use of generative AI for content development. In that setting, the relevant question is not whether the model is authoritative, but whether it can speed up drafting while leaving judgment, verification, and adaptation with the instructor [13, 15].

Generative AI also creates familiar risks. Models may produce inaccurate statements with a confident tone, cite non-existent or mismatched sources, or drift into a register that does not fit the intended audience. For that reason, this study treated model outputs as drafts requiring source checks, disciplinary review, and student-facing readability edits rather than as finished instructional material [13, 15].

3 Methodology

This study draws on constructivist learning theory as its primary analytical framework. Constructivism, rooted in the work of Piaget, Vygotsky [16], and subsequent scholars [8], holds that knowledge is actively constructed by learners through their interactions with the environment, with other learners, and with culturally situated tools and artifacts. Kolb's experiential learning cycle [22] posits that concrete experience, reflective observation, abstract conceptualization, and active experimentation. As a result, it provides a useful heuristic for understanding how scenario-based activities can structure learning.

Within this framework, the scenarios function as mediating artifacts: they present challenges that are sufficiently complex to require guided effort but sufficiently structured to remain tractable for the target learner population. The human–AI workflow was also iterative: faculty and student reviewers refined draft material through cycles of generation, evaluation, and revision.

Connectivism, as proposed by Siemens [24], further informs the study by highlighting how learning in networked environments depends on the learner's capacity to identify, evaluate, and synthesize information drawn from diverse sources, a capacity that scenario-based exercises are well positioned to develop.

3.1 Course Context and Design

The pilot was conducted as part of an undergraduate course in privacy, ethics, and policy in information security offered in the cybersecurity program at Fairmont State University. The course satisfies requirements aligned with the national cybersecurity workforce framework knowledge units [4] and enrolls students in their junior or senior year. The course is organized into six instructional modules:

- *Module 1: Foundations of applied ethics.* Classical ethical theories (deontology, consequentialism, virtue ethics) as they apply to digital environments and data stewardship.

- *Module 2: Privacy regulation and compliance.* Examination of GDPR, CCPA, HIPAA, and sector-specific privacy frameworks, with emphasis on jurisdictional complexity.
- *Module 3: Organizational data governance.* Policies, procedures, and ethical obligations surrounding data collection, retention, sharing, and disposal within corporate settings.
- *Module 4: Incident response and disclosure ethics.* The ethical dimensions of breach notification, vulnerability disclosure, and organizational accountability.
- *Module 5: Surveillance, civil liberties, and national security.* Balancing state security interests against individual privacy rights, with attention to mass surveillance programs and lawful interception.
- *Module 6: Emerging challenges.* AI governance, algorithmic fairness, biometric data, and the ethical trajectory of data-intensive technologies.

Our scenario development process followed an iterative human–AI workflow. The faculty lead (KW) authored detailed specification documents for each module, including the learning objectives, required conceptual coverage, target complexity level, desired document types (e.g., regulatory correspondence, internal emails, meeting transcripts, news articles), and constraints on length and tone. These specification documents were provided to a commercial LLM as contextual input alongside excerpts from the course lecture materials and selected readings.

The LLM produced initial drafts, and the research team reviewed them before any material reached students. Three student collaborators, that is, two graduate research assistants (MR and AC) and one advanced undergraduate (TO), provided feedback on readability, engagement, and disciplinary accuracy. The graduate assistants had previously completed the course and brought firsthand experience of the instructional context. Feedback was consolidated and used to revise the prompt specifications; revised prompts were then resubmitted to the model. Each module went through three or four draft-review cycles before the team accepted a classroom version.

Each finalized scenario adopted a dossier format containing the following components:

- A narrative overview establishing the situational context and identifying the central ethical dilemma.
- Background documents providing regulatory, organizational, and technical context.
- Character profiles of key stakeholders, each representing a distinct perspective (e.g., corporate executive, privacy officer, affected consumer, regulatory investigator).
- A chronological timeline of events leading to the dilemma.
- Supplementary artifacts: simulated emails, press releases, regulatory notices, social media posts, and internal memoranda.

For example, Module 3 featured a scenario in which a mid-sized health technology firm discovered that a third-party analytics vendor had been aggregating de-identified patient data in ways that enabled re-identification. The dossier included a simulated letter from a state attorney general's office, internal Slack messages between engineering and legal teams, a draft press statement, and excerpts from the vendor's data processing agreement. Students were asked to analyze the obligations of each stakeholder, evaluate the adequacy of the firm's data governance framework, and propose remedial actions consistent with applicable regulations.

3.2 Teaching Integration

A structured teaching guide accompanied each scenario module. The guide followed the following instructional sequence, also shown in Figure 1.

1. *Contextualization.* The instructor situates the scenario within the module's learning objectives and provides a brief background on the relevant regulatory or ethical framework.
2. *Independent analysis.* Students read the dossier individually and complete a structured response worksheet identifying key stakeholders, competing interests, applicable legal standards, and preliminary ethical assessments.
3. *Small-group deliberation.* Students discuss their analyses in groups of three to four, with each group tasked to reach a consensus recommendation.
4. *Full-class synthesis.* Groups present their recommendations; the instructor facilitates comparison across groups and draws connections to theoretical frameworks.
5. *Reflective writing.* Students submit a brief individual reflection on how the exercise affected their understanding of the module topic.



Figure 1: The instructional sequence adopted in our work.

Table 1: Percentage of students who agreed or strongly agreed with statements about the AI-assisted learning scenarios.

Statement	M	SD	% Agree
The scenarios captured my interest.	4.57	0.51	95.2
The scenarios helped me understand key course concepts.	4.33	0.73	90.5
The scenarios were more engaging than traditional lectures.	4.29	0.78	90.5
The scenarios motivated me to think carefully about the issues.	4.38	0.67	95.2
The scenarios increased my confidence in applying course concepts.	4.19	0.75	85.7
The scenarios helped me appreciate the complexity of the subject.	4.24	0.70	90.5

Note. Agreement reflects the percentage of students who selected “Agree” or “Strongly Agree” on a 5-point Likert scale (1 = Strongly Disagree, 5 = Strongly Agree).

This sequence reflects the experiential learning cycle described by Kolb [22]: the scenario provides a concrete experience; independent and group analyses encourage reflective observation and abstract conceptualization; and the reflective writing assignment invites active experimentation by applying new understandings to broader contexts.

3.3 Evaluation Protocol

The evaluation used a mixed-methods protocol to examine students’ perceptions of the scenarios. The protocol comprised quantitative and qualitative components. A post-intervention survey administered during the final week of the semester included six five-point Likert scale items [25] targeting three constructs: student engagement with the scenario materials, perceived learning outcomes, and perceived effectiveness of the scenarios in addressing ethical, legal, and policy content. Response anchors ranged from 1 (Strongly Disagree) to 5 (Strongly Agree). Three open-ended items invited students to describe (a) aspects of the scenarios they found most valuable for their learning, (b) challenges or limitations they encountered, and (c) suggestions for improvement. Responses were analyzed using Braun and Clarke’s [26] six-phase thematic analysis procedure.

All enrolled students were invited to participate. Responses were collected anonymously; no compensation or course credit was offered for participation. Of 23 enrolled students, 22 submitted surveys (return rate: 95.7%). One response was excluded because of incomplete data, yielding an analytic sample of 21 students (91.3% of enrolled students).

The study was reviewed and approved by the university’s institutional review board. All participants provided informed consent. The use of self-reported perceptual data represents a recognized limitation; strategies for triangulating these data with performance-based measures are addressed in the Discussion.

4 Results

The Likert items showed positive perceptions across the measured constructs. Table 1 reports means, standard deviations, and the percentage of respondents selecting “Agree” or “Strongly Agree” for each item.

The highest-rated item was interest capture (4.57 ± 0.51), with 95.2% of respondents agreeing or strongly agreeing. Items related to perceived understanding of key concepts (4.33 ± 0.73) and motivation for careful thinking (4.38 ± 0.67) also received favorable ratings above 90%. The item measuring confidence in applying course concepts received the lowest agreement percentage (85.7%), though the mean rating (4.19) remained well above the scale midpoint.

Table 2 presents the perceived effectiveness of the scenarios in addressing each of the three core content domains.

Table 2: Percentage of students who rated the scenarios as “Extremely effective,” “Very effective,” or “Moderately effective” in addressing each content domain ($n = 21$).

Content Domain	% Effective
Ethics	95.2
Privacy regulation	90.5
Organizational policy	85.7

Note. Effectiveness ratings reflect the percentage of students who selected “Extremely effective,” “Very effective,” or “Moderately effective” on a 5-point scale.

Ethics content received the highest effectiveness rating (95.2%), followed by privacy regulation (90.5%) and organizational policy (85.7%). This pattern is consistent with the idea that dossier-style scenarios are especially useful for ethics content, where a narrative can make competing values visible. The survey design does not show why the ethics item was rated highest, so this interpretation should be treated as tentative.

The open-ended responses clustered into four themes. Theme 1 involved realism and contemporary relevance. Students often pointed to realism as a reason the exercises held their attention. One respondent wrote: “The scenario about the health data re-identification felt like something that could actually happen at a company I might work for and it made me take the exercise seriously.” Another noted that “reading simulated emails between a legal team and engineers made the dilemma feel concrete rather than abstract.” The inclusion of contemporary regulatory references (e.g., GDPR articles, CCPA provisions) was cited by multiple respondents as a feature that connected classroom learning to professional practice.

The second theme involved the different stakeholder perspectives. Several respondents highlighted the value of encountering multiple viewpoints within a single scenario. One student commented: “Having to think about what the privacy officer would do versus what the CEO would want forced me to see the tradeoffs instead of jumping to one answer.” This finding aligns with the constructivist emphasis on perspective-taking as a mechanism for deepening understanding [8].

The third theme involved vocabulary and complexity challenges. A subset of respondents ($n = 4$) noted that certain scenarios contained legal or technical terminology that exceeded their current knowledge level. One wrote: “Some of the regulatory language was dense. I had to look up several terms before I could fully follow the scenario.” In response, the research team developed supplemental glossaries for each module. Informal comments after the glossaries were introduced suggested that they reduced, but did not eliminate, comprehension barriers.

Finally, the last theme regarded the desire for expanded interactivity. Three respondents expressed a desire for scenarios that incorporated more interactive elements, such as branching decision paths or simulated stakeholder negotiations. While the current scenarios were designed as static dossiers, this feedback points toward future design iterations that could integrate interactive or game-based elements [23].

5 Discussion

The survey results suggest that the AI-assisted workflow produced materials that students in this course perceived as engaging, relevant, and useful for learning about data privacy and security ethics. Agreement percentages ranged from 85.7% to 95.2% across the six Likert items, which indicates broad endorsement within this small class.

The strongest ratings appeared on items related to interest capture and motivation for careful thinking. This pattern is consistent with the theoretical expectation that scenario-based learning, by embedding abstract concepts within concrete narrative contexts, activates experiential engagement in ways that traditional lecture formats do not [7,22]. The lower (though still positive) ratings on confidence in applying concepts suggest that while the scenarios effectively stimulated interest and analytical engagement, additional scaffolding may be needed to support the transfer of scenario-based learning to novel professional situations. This transfer gap is well documented in the constructivist literature and may be addressed through cumulative portfolio activities or capstone exercises that require students to apply ethical reasoning across multiple contexts [17].

The domain-specific effectiveness ratings further suggest that the scenario format is particularly suited to ethics instruction, where narrative dilemmas naturally foreground competing values and require evaluative judgment. Privacy regulation and organizational policy, which involve more procedural and factual content, may benefit from supplementary instructional strategies that complement the scenario exercises.

The qualitative responses help explain the survey pattern. The prominence of realism and contemporary relevance as themes suggests that scenario fidelity mattered for engagement in this course. This finding is consistent with prior work on case-study pedagogies [11] and situated learning theory [17]. Students did not merely acknowledge the realism of the scenarios in general terms; they cited specific artifacts (simulated emails, regulatory letters) that anchored abstract concepts in tangible professional contexts. This specificity suggests that the dossier format contributed to perceived authenticity.

The theme regarding multiple stakeholders is also important pedagogically. By presenting dilemmas through the eyes of multiple actors with divergent interests, the scenarios required students to engage in perspective-taking, a cognitive process that constructivist theory identifies as central to knowledge construction [16]. Several students described this process as challenging but ultimately valuable, noting that it disrupted their tendency to apply a single ethical framework to every situation.

The vocabulary challenge identified by a minority of respondents needs attention in future iterations. This finding echoes broader concerns about the register and lexical choices of LLM-generated text, which may default to a level of formality or specialization that exceeds the expectations of introductory-level students [13]. The glossary appeared useful in the present course, but its effect on comprehension should be measured directly in future work.

The workflow offers a practical model for using generative AI in scenario development without treating the model as an authority. Three features mattered in practice. First, the specification documents provided to the LLM served as detailed creative briefs, constraining the model's output across content, format, tone, and complexity. This structured prompting approach reduced the incidence of off-topic or inappropriately pitched content. Second, the inclusion of student collaborators in the review cycle introduced perspectives that faculty alone might overlook. This is particularly true regarding readability, engagement, and the experiential resonance of specific narrative details.

Third, iterative revision created a safeguard against factual inaccuracies, invented citations, and mismatched regulatory references. By treating each AI-generated draft as raw material subject to expert verification, the workflow positioned the LLM as a productivity tool rather than an authoritative source. This framing matters: generative models can produce plausible but incorrect statements, and the consequences of embedding such inaccuracies in educational materials are nontrivial, especially in a domain where legal and regulatory precision matters [15].

Because this study is still in progress, we emphasize the provisional nature of some elements of our work. Although this paper uses the pilot results to motivate and refine the research program rather than to make definitive causal claims about the superiority of AI-assisted scenario design, the study suggests several practical implications for information security educators.

Scalable content development. The human–AI collaborative workflow reduced the time required to develop each scenario module from an estimated 15–20 hours (based on the lead author's prior experience with manual scenario construction) to approximately 5–7 hours, inclusive of specification authoring, AI interaction, and iterative revision. For programs updating scenario libraries across multiple courses, that difference may be consequential, although it is based on a single instructor's prior experience rather than a controlled time study.

Customization for diverse student populations. Because the LLM can generate scenario variants at relatively low marginal cost, the workflow supports the production of culturally, sectorally, or regulatorily varied versions of a single underlying scenario. This adaptability has implications for programs that serve diverse student populations or that wish to align content with regional regulatory frameworks.

Faculty development. Adopting AI-assisted content development requires faculty to develop new competencies in prompt engineering, AI output evaluation, and iterative design. Faculty adopting this workflow would benefit from shared review checklists, examples of prompt specifications, and peer review of AI-assisted materials.

6 Conclusion

This pilot study examined the use of a generative AI model as a drafting partner in developing scenario-based instructional materials for an undergraduate data privacy and security ethics course. A structured iterative workflow, grounded in constructivist principles, produced dossier-format scenarios that integrated multiple document types, stakeholder perspectives, and contemporary regulatory contexts. Pilot evaluation results from 21 students indicated high levels of perceived engagement, understanding, and scenario effectiveness, with the strongest outcomes observed for ethics content. Qualitative analysis confirmed the importance of realism, multi-perspective design, and contemporary relevance in driving student engagement, while also identifying vocabulary complexity as a challenge amenable to targeted intervention.

The findings suggest that human–AI collaborative content development can be a workable approach for producing high-quality, contextualized learning materials in the information security domain. At the same time, we acknowledge the limitations of the present study, including the small sample size, self-reported outcomes, and the absence of a comparison condition. Future work should test this approach against instructor-authored scenarios and measure learning outcomes directly, so that any efficiency gains are evaluated alongside instructional quality.

The next step is a stronger evaluation. A controlled experimental study comparing student outcomes (engagement, knowledge acquisition, ethical reasoning proficiency) between AI-assisted and manually developed scenarios would establish the relative contribution of the AI component. Longitudinal tracking of students who have completed scenario-based modules could assess whether perceived learning gains translate into durable changes in ethical awareness and professional behavior. Investigations into how students distinguish between AI-generated and human-authored instructional materials, and whether this distinction influences their engagement,

would inform transparency practices in AI-assisted pedagogy. These future-research directions also identify the main work-in-progress elements of the project. In particular, we continue to revise the scenario library, develop performance-based assessment rubrics, and design a stronger comparative study to distinguish the effects of scenario-based pedagogy from those of AI-assisted content generation.

References

- [1] R. von Solms and J. van Niekerk, "From information security to cyber security," *Computers & Security*, vol. 38, pp. 97–102, 2013.
- [2] ISC2, "Cybersecurity workforce study." <https://www.isc2.org/Research/Workforce-Study>, 2022. Accessed 2026-05-14.
- [3] W. A. Conklin, R. E. Cline, and T. Roosa, "Re-engineering cybersecurity education in the US: An analysis of the critical factors," in *Proceedings of the 47th Hawaii International Conference on System Sciences*, pp. 2006–2014, 2014.
- [4] W. Newhouse, S. Keith, B. Scribner, and G. Witte, "NICE cybersecurity workforce framework," Special Publication 800-181, National Institute of Standards and Technology, Gaithersburg, MD, 2017.
- [5] A. Parrish, J. Impagliazzo, R. K. Raj, H. Santos, M. R. Asghar, A. Jøsang, T. Pereira, and E. Stavrou, "Global perspectives on cybersecurity education for 2030: A case for a meta-discipline," in *Proceedings Companion of the 23rd Annual ACM Conference on Innovation and Technology in Computer Science Education*, pp. 36–54, 2018.
- [6] V. Švábenský, J. Vykopal, and P. Čeleda, "What are cybersecurity education papers about? A systematic literature review of SIGCSE and ITICSE conferences," in *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*, pp. 2–8, 2020.
- [7] C. E. Hmelo-Silver, "Problem-based learning: What and how do students learn?," *Educational Psychology Review*, vol. 16, no. 3, pp. 235–266, 2004.
- [8] D. H. Jonassen, "Designing constructivist learning environments," in *Instructional-Design Theories and Models: A New Paradigm of Instructional Theory* (C. M. Reigeluth, ed.), vol. 2, pp. 215–239, Mahwah, NJ: Lawrence Erlbaum Associates, 1999.
- [9] R. Beuran, D. Tang, C. Pham, K. ichi Chinen, Y. Tan, and Y. Shinoda, "Integrated framework for hands-on cybersecurity training: CyTrONE," *Computers & Security*, vol. 78, pp. 43–59, 2018.
- [10] R. S. Weiss, F. Turbak, E. Mache, and M. E. Locasto, "Teaching cybersecurity analysis skills in the cloud," in *Proceedings of the 46th ACM Technical Symposium on Computer Science Education*, pp. 332–337, 2015.
- [11] W. He, L. Xu, G. Kruck, and D. S. Comer, "Teaching information security with workflow technology," *Journal of Information Systems Education*, vol. 25, no. 3, pp. 201–210, 2014.
- [12] N. Chowdhury and V. Gkioulos, "Cyber security training for critical infrastructure protection: A literature review," *Computer Science Review*, vol. 40, p. 100361, 2021.
- [13] E. Kasneci, K. Seßler, S. Küchemann, M. Bannert, D. Dementieva, F. Fischer, U. Gasser, G. Groh, S. Günne-mann, E. Hüllermeier, S. Krusche, G. Kutyniok, T. Michaeli, C. Nerdel, J. Pfeffer, O. Poquet, M. Sailer, A. Schmidt, T. Seidel, M. Stadler, J. Weller, J. Kuhn, and G. Kasneci, "ChatGPT for good? On opportunities and challenges of large language models for education," *Learning and Individual Differences*, vol. 103, p. 102274, 2023.
- [14] J. Zhu, C. Yang, and T. Wang, "ChatGPT and education: A comprehensive review," *Education and Information Technologies*, 2023.
- [15] T. Farrelly and N. Baker, "Generative artificial intelligence: Implications and considerations for higher education practice," *Education Sciences*, vol. 13, no. 11, p. 1109, 2023.
- [16] L. S. Vygotsky, *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA: Harvard University Press, 1978.
- [17] J. Lave and E. Wenger, *Situated Learning: Legitimate Peripheral Participation*. Cambridge: Cambridge University Press, 1991.

- [18] M. Dark, J. Mirkovic, R. Liles, and J. Coffman, "Assessing student performance and curricula efficacy," *IEEE Security & Privacy*, vol. 13, no. 4, pp. 84–87, 2015.
- [19] H. T. Tavani, *Ethics and Technology: Controversies, Questions, and Strategies for Ethical Computing*. Hoboken, NJ: John Wiley & Sons, 4 ed., 2013.
- [20] M. J. Quinn, *Ethics for the Information Age*. New York: Pearson, 8 ed., 2020.
- [21] E. A. Buchanan, "Internet research ethics: Past, present, and future," in *The Handbook of Internet Studies* (M. Consalvo and C. Ess, eds.), pp. 83–108, Oxford: Wiley-Blackwell, 2011.
- [22] D. A. Kolb, *Experiential Learning: Experience as the Source of Learning and Development*. Englewood Cliffs, NJ: Prentice-Hall, 1984.
- [23] S. Deterding, D. Dixon, R. Khaled, and L. Nacke, "From game design elements to gamefulness: Defining "gamification"," in *Proceedings of the 15th International Academic MindTrek Conference*, pp. 9–15, 2011.
- [24] G. Siemens, "Connectivism: A learning theory for the digital age," *International Journal of Instructional Technology and Distance Learning*, vol. 2, no. 1, pp. 3–10, 2005.
- [25] R. Likert, "A technique for the measurement of attitudes," *Archives of Psychology*, vol. 22, no. 140, pp. 5–55, 1932.
- [26] V. Braun and V. Clarke, "Using thematic analysis in psychology," *Qualitative Research in Psychology*, vol. 3, no. 2, pp. 77–101, 2006.