Paddling Without a Map: The Understructured Rise of AI Literacy in Internet-Enabled Higher Education

Heather Horowitz (ORCID: https://orcid.org/0009-0000-6640-7485), SUNY Empire State; Jeff Foulkes (ORCID: https://orcid.org/0009-0003-6179-3884), Glendale Community College; Pam Doran (ORCID: https://orcid.org/0009-0003-6647-8007), SUNY Empire State, USA

Abstract

Higher education institutions are rapidly implementing AI literacy initiatives through internet-enabled courses, workshops, and professional development experiences, yet these efforts are often undertaken without sustained engagement with empirically grounded AI literacy frameworks. Despite the availability of such frameworks, adoption across institutions remains limited, pointing to gaps in dissemination, perceived relevance, and institutional prioritization. Drawing on publicly available materials and semi-structured interviews across 17 higher education institutions, this qualitative study examines how institutions govern and design AI literacy learning experiences for non-technical learners under conditions of rapid technological change. Findings indicate that most institutions rely on informal influences, independent research, and internally defined priorities rather than formal frameworks when designing internet-enabled AI literacy offerings. Institutions primarily emphasized functional use and ethical caution, with less consistent attention to social, civic, and global dimensions of AI. While this decentralized approach affords flexibility and rapid response to technological change, it raises concerns about coherence, equity, and consistency in AI literacy implementation. The study also identifies notable misalignment between publicly stated workforce-readiness goals and enacted instructional design, as well as tensions between tool-agnostic approaches and the implementation of accessibility and Universal Design for Learning (UDL) practices. By providing a cross-institutional empirical account of current AI literacy practices, this study contributes to scholarship on internet-enabled teaching and learning in higher education and offers implications for institutional policy, professional development strategy, and the development of more coherent and equitable AI literacy initiatives.

Keywords: AI in Higher Education, AI Literacy, Institutional Adoption, AI Frameworks, AI Learning Outcomes

Introduction

Artificial intelligence (AI) is reshaping educational practice and workforce demands, creating new pressure on higher education institutions to prepare learners for an AI-driven world (Daher, 2025). As AI technologies become increasingly integrated into the workforce and society, developing student AI literacy has emerged as an institutional imperative (Almatrafi et al., 2024; Southworth et al., 2023).

Yet the rapid acceleration of generative AI since 2022 has pushed institutions to respond quickly, often before shared definitions, learning outcomes, or design frameworks have solidified. While many institutions are rapidly implementing AI literacy initiatives, little is known about the intended learning outcomes or the frameworks guiding these experiences, critical factors for ensuring coherent, inclusive, and effective design.

This rapid uptake positions AI literacy as a technology-in-practice adoption challenge: institutions are implementing AI literacy under substantial time pressure, typically without shared frameworks to guide decisions about learning outcomes, accessibility, tool use, or ethical considerations. As a result, many AI literacy offerings are emerging through localized improvisation rather than coordinated institutional strategy, producing variability in coherence and equity. Because AI literacy initiatives are increasingly delivered through online, hybrid, and digitally mediated formats, these patterns reflect broader, well-documented dynamics in educational technology adoption within higher-education contexts.

This study addresses this gap in structure and implementation by examining how higher education institutions design AI literacy learning experiences for non-technical learners, analyzing both publicly available materials and qualitative interviews to identify design patterns, guiding frameworks, institutional strategies, and misalignment. The following research questions guided this study:

1. What, if any, AI literacy frameworks guide the design of AI learning experiences?
2. What are the AI learning experience outcomes across institutions?

This study makes three contributions to the field of educational technology. First, we analyze 17 institutional AI literacy courses and nine educator interviews to provide a cross-institutional empirical account of how higher education institutions are designing AI literacy for non-technical learners. Second, we identify consistent patterns across institutions, including minimal framework adoption, reliance on local “AI champions”, and key misalignments between public claims (e.g., workforce readiness, tool-neutrality) and actual curricular design, including implications for universal design for learning (UDL) and accessibility. Third, we demonstrate that AI literacy development is functioning as a decentralized ed-tech adoption process, and offer insights to support more coherent, equitable, and intentional AI literacy implementation.

To guide this study, we draw on a pragmatic, socio-technical perspective that views AI literacy as shaped through the interaction of institutional priorities, instructional design decisions, and rapidly evolving technological contexts. This perspective aligns with prior work emphasizing that AI integration in education is not solely a technical challenge but also an organizational and pedagogical one, requiring institutions to balance innovation with coherence, equity, and accessibility (Kassorla et al., 2024; Robert, 2025). Within this framing, AI literacy is understood not as a fixed set of competencies, but as a context-dependent construct influenced by local decision-making, available resources, and external pressures for workforce alignment. This lens allows the study to examine how institutions interpret, adapt, and operationalize AI literacy across internet-enabled learning environments, and how these choices shape the alignment between stated goals and enacted instructional practices.

Literature Review

Defining AI Literacy in Education

Artificial intelligence (AI) influences day-to-day academic work far more than it did even a few years ago. Many instructors have adjusted course design, grading practices, and support strategies in response to the emergence of these tools. Over the past 20 years, work in artificial intelligence in education (AIEd) has grown in uneven but noticeable ways. Recent analyses of AIEd research trends show accelerated growth after 2012, which created disparate levels of adoption across disciplines. Different fields have picked up pieces of the conversation as their needs and interests emerged (Akhmadieva et al., 2024). According to Kennedy and Gupta (n.d.), current efforts to promote AI literacy in higher education focus on data-driven personalization, intelligent tutoring systems, natural language processing, and adaptive learning environments. The exponential growth of AI tools since ChatGPT's launch in 2022 has put significant pressure on both learners and educators. There is a need for clearer explanations of how these systems work, how their outputs should be interpreted, and how to navigate the ethical and social questions they raise. In response, many educators now treat AI literacy as a central institutional goal as it gives students a practical way to engage with these technologies responsibly (Kassorla et al., 2024; Southworth et al., 2023).

Early work in this area treats AI literacy less as a narrow technical skill and more as a collection of practical competencies that students can apply across settings. Long and Magerko (2020) define AI literacy as the capacity to understand, evaluate, and use AI tools in everyday contexts. They focus particularly on learners without computing backgrounds. Their work shifts focus away from coding tasks to instead examine broader habits of thinking, including conceptual understanding, critical questioning, and ethical judgment. This view suggests that non-technical learners benefit from recognizing AI systems when they encounter them, and helps them think carefully about these systems' limits and reflect on the social impact of their use.

Building on this line of inquiry, Ng et al. (2021) describe AI literacy as a broad, layered set of abilities. They point to understanding core ideas, such as how AI systems learn from data, generate outputs, and operate within social contexts, and using AI tools, evaluating them, and creating systems while engaging with the ethical questions that come from use of these tools. Their review highlights the lack of shared definitions and assessment approaches across educational contexts. Although their primary focus is K–12 education, they also note that higher education runs into many of the same problems: institutions often introduce AI tools before anyone has agreed on the learning goals they are meant to support. Taken together, these definitions show AI literacy as a broad educational construct that blends conceptual understanding, hands-on use, and thoughtful engagement with ethical questions, but one that may lack alignment.

AI Literacy Frameworks and Competency Models

As interest in AI has increased, educators and researchers have created multiple frameworks to guide curriculum planning and learning outcomes. Several emphasize progression across levels of engagement with AI systems. The AI4K12 Initiative lays out five core ideas: Perception, Representation and Reasoning, Learning, Natural Interaction, and Societal Impact. These ideas serve as starting points for teaching AI concepts in a clear and approachable way (Touretzky et al., 2022). Although created with K–12 settings in mind, the AI4K12 approach has started to shape work within higher education. Some colleges draw on this structure when developing introductory AI experiences for students who are not majoring in technical fields.

Some frameworks place stronger emphasis on ethical engagement and critical evaluation. Mills et al. (2024) offer a model that highlights understanding, use, and evaluation of AI. They argue that thoughtful evaluation and ethical reflection should shape how people decide when and why to use these tools. Their stance aligns with higher education goals, where faculty often emphasize critical thinking and responsible decision-making when integrating new technologies.

More comprehensive competency models have appeared at global and institutional levels. UNESCO's AI Competency Framework for Students (2024) outlines competencies across four dimensions: human-centered mindset, ethics of AI, AI techniques and applications, and AI system design. The framework connects AI literacy to social responsibility and sustainable development.

Within U.S. higher education, EDUCAUSE's AI Literacy in Teaching and Learning framework takes a role-based approach. This structure reflects a growing recognition that AI literacy needs differ among students, faculty, and staff, which is particularly relevant in higher-education settings where institutional responsibilities vary widely. It breaks competencies into what students, faculty, and staff each need to understand (Kassorla et al., 2024). Kennedy and Gupta (n.d.) offer a more blended structure that ties together technical skills, cognitive processes, ethical reasoning, and the social contexts in which AI tools are used.

Institutional case studies offer a closer look at how colleges and universities adapt these ideas to fit their own priorities and constraints. Hibbert et al. (2024) outline Barnard College's four-level AI literacy pyramid: Understand, Use, Analyze/Evaluate, and Create. Horowitz and Foulkes (2025) extend this line of thinking through the Inclusive, Critical, and Applied AI Literacy (ICAIL) framework, which draws on common AI literacy dimensions such as understanding, ethical considerations, and use, while expanding beyond them to center social responsibility, digital citizenship, and equity-oriented considerations. Their model places non-technical learners at the core and incorporates elements of critical pedagogy and Universal Design for Learning to support a wider range of learning needs. These examples show how individual institutions are designing structured approaches, but they also reveal inconsistencies within the broader field.

Recent work highlights assessment tools such as the GenAI Literacy Assessment Test (GLAT), a 20-item instrument that measures generative AI literacy (Jin et al., 2024). Jin et al. (2024) found that many learners can make substantial progress in AI literacy without advanced programming knowledge. The authors show that accessible, layered approaches often produce stronger results because they meet students where they are.

Persistent Gaps and the Need for Empirical Examination

The literature shows broad agreement that AI literacy matters, yet the way it is implemented across higher education varies. Valentini and Blancas (2025) note that many institutions rely on reactive measures rather than coherent, long-term planning. They also note the lack of higher-education-specific competency frameworks to bring greater consistency to this work. Without shared language or clear outcomes, AI literacy initiatives risk becoming uneven across student populations.

The literature identifies several persistent gaps that merit further examination. These include unclear learning outcomes, uneven use of existing frameworks, and limited focus on how AI literacy experiences are shaped for learners without technical backgrounds. Scholars have outlined possible definitions and goals for AI literacy, yet far less is known about how colleges and universities operationalize those ideas in practice. To move the field forward, researchers need more empirical work that clarifies design choices, surfaces the influences shaping them, and shows where common patterns are beginning to emerge (Chiu et al., 2024).

This study examines how higher education institutions design AI literacy learning experiences for non-technical learners. Instead of proposing a new framework, the study investigates existing institutional practices to determine guiding learning outcomes, use of formal frameworks, and operationalization across contexts. Informed by The Evolution of Research on AI and Education Across Four Decades, the authors use a scaffolded framework to examine publicly available resources and conduct interviews with instructional leads, aiming to provide insights into how current approaches foster coherence, inclusivity, and intentional design in AI literacy within higher education (Rismanchian & Doroudi, 2025). This structured approach informs the research questions guiding the present study.

Methodology

Online Material (Data A)

The research team identified AI literacy learning experiences intended for non-technical learners at higher education institutions (n = 35). Following an initial review of publicly available website materials, 17 experiences were selected for inclusion based on the following criteria:

1. The AI learning experience was not a full academic program.
2. The AI learning experience was created by a higher education institution.
3. The target audience was non-technical learners.
4. While modality was not restricted, online experiences were required to include at least one synchronous component.

One institution was subsequently removed because the advertised experience was not affiliated with the institution. Upon closer review, the offering could not be verified as an official institutional initiative. As a result, it was excluded to maintain the integrity of the dataset.

Data Collection

To examine the design of AI literacy learning experiences, we conducted a systematic review of publicly available materials from higher education institutions. This review aimed to capture key elements including program structure, purpose, stated learning outcomes, and references to guiding frameworks. Data collection was guided by a structured review sheet containing prompts to ensure consistent documentation across institutions. Institutions were anonymized and labeled A–Q.

Upon further review, two institutions were removed from the dataset. Institution D did not meet the inclusion criteria, and Institution P was excluded after follow-up verification revealed that the listed AI course was not institutionally affiliated and therefore did not represent an institutional AI literacy offering.

Researchers independently examined course websites and related materials, recording detailed notes in the review sheet. Independent review was emphasized to maintain analytic rigor and avoid cross-contamination of observations. Completed review sheets were uploaded to a shared repository for subsequent coding and thematic analysis.

Following the initial review, the research team conducted a second round of data collection. Researcher 1 reviewed all selected AI literacy learning experiences across both rounds to ensure consistency. A second researcher independently reviewed a subset of institutions; importantly, no institution was reviewed twice by the same secondary reviewer, preserving independence of observations.

Online Material (Data A) Analysis

The research team employed a hybrid coding approach aligned with a pragmatic qualitative research paradigm (Patton, 2015), integrating both inductive and deductive reasoning. Researcher 1 developed an initial codebook with definitions and examples, which was subsequently refined and validated by Researchers 2 and 3.

Coding began inductively, with Researcher 1 reviewing a subset of institutions and generating preliminary codes based on emergent patterns. Digital analytic tools were used to support early code organization, after which the content was synthesized into a structured codebook. The coding process remained flexible to allow for the addition of emergent codes (Deterding & Waters, 2021).

To ensure consistency in code application, Researcher 1 trained the research team using the codebook. The team piloted three institutions together, comparing and discussing coding decisions. Two additional institutions were then coded independently, followed by a group discussion to resolve discrepancies. After calibration, two researchers coded the remaining institutions.

Following initial coding, the research team engaged in iterative analytic discussions to address ambiguous excerpts and refine code definitions. During this phase, the code *Strategic_Planning* was redefined as *Strategic_Vision*, and new codes (*Workforce_Ready* and *Resources*) were introduced. Coding sheets were updated to reflect these refinements.

Semantic Jaccard similarity scores were used as a diagnostic tool to assess coding consistency. The Jaccard similarity coefficient measures the degree of overlap in coding between researchers, with higher values indicating greater agreement. Initial overall similarity was 0.57. Researcher 1 reviewed similarity outputs to identify areas requiring further clarification and met individually with each researcher to align code interpretation.

Codes with the highest discrepancies included: Other, UDL_Engagement, Learning_Outcomes, Resources, AI_Centers_or_Staff, Support_Diverse_Learners, Workforce_Ready, and Non_Technical. After refinement and discussion, similarity scores were recalculated, resulting in an overall score of 0.69, indicating improved shared understanding across the research team. These scores were used to guide collaborative interpretation and codebook refinement rather than as fixed thresholds for reliability. A breakdown of similarity scores by institution is presented in Table 1.

Table 1
Jaccard Similarity Score

Institution	Similarity Score
A	.75
B	.53
C	.56
E	.58
F	.58
G	.71
H	.92
I	.38
J	.50
K	.40
L	.47
M	.46
N	.53
O	.50
Q	.67

Interviews (Data B)

After compiling contact information for 16 institutions and receiving IRB approval from SUNY Empire State, Researcher 1 sent direct email invitations to potential participants, inviting them to participate in semi-structured interviews focused on the design and implementation of their AI literacy learning experiences. Nine participants agreed to be interviewed. Interviews were distributed among the research team based on availability.

To support triangulation between Data A and Data B, the same A–Q institutional labels used in the document review were retained for interview data. Each interviewer coded their own transcripts using the shared codebook, ensuring consistency in analytic approach across data sources.

Data Synthesis

To ensure consistency in data extraction and interpretation, Researcher 1 developed a synthesis template aligned with each research question. This template guided how researchers integrated information from both Data A and Data B, supporting standardized data capture, facilitating thematic coding, enabling triangulation across data sources, and informing the interpretation of institutional patterns and implications.

Findings

Guiding AI Literacy Frameworks

Only one institution (A) explicitly referenced an established AI literacy framework, AI4K12.org's *Big Five AI Ideas*. At the time of data collection, two institutions (N and O) were in the process of developing or had established institutional frameworks to guide their AI literacy learning experiences. Institution N noted that "there is not a single AI framework" and described efforts to define AI literacy in relation to multiple existing information literacy frameworks. Institution O developed a set of AI core competencies that serve as the foundation for an AI microcredential, with the reviewed course embedded within that broader structure.

Interview data provide additional insight into why institutions operate outside formal frameworks. Participants described design decisions grounded in independent research (M, J) or the adaptation of non-AI-specific frameworks, approaches that are not typically visible in public-facing materials. For example, Institution M described assembling its approach by integrating elements from multiple sources. Other institutions drew on peer models: Institution J consulted with a peer institution serving a similar student population, while Institution N adapted components from another institution's course, without clear alignment to a single AI literacy framework.

Across institutions, most AI learning experiences were developed using human-centered approaches, backward design principles, or informal influences such as the ACM curriculum and accessibility frameworks, rather than formal AI literacy frameworks. Institution B, for instance, reported limited awareness of AI-specific frameworks and noted that the ACM computer science curriculum indirectly informed its offering: "We used the ACM list in our computer science course, which kind of informed the Elements of AI [experience]." Collectively, these findings indicate that while institutions are actively designing AI literacy experiences, they do so largely without reliance on formally established frameworks.

Institutional Learning Outcomes

Analysis of institutional documents (Data A) and interview transcripts (Data B) revealed that learning outcomes cluster into three primary domains: Technical (LO_Tech), Ethical (LO_Ethical), and Social (LO_Social). While technical and ethical outcomes were consistently present across institutions, social outcomes and workforce readiness showed notable variation between publicly stated goals and actual instructional implementation.

Technical literacy emerged as the most prominent domain, appearing in 14 sources. Rather than emphasizing advanced technical processes such as coding or model development, these outcomes focused on foundational AI concepts appropriate for non-technical learners. Common objectives included prompt engineering and large language model (LLM) fluency, understanding how generative AI operates (e.g., probability and prediction), and distinguishing between different types of machine learning (e.g., supervised versus unsupervised learning). Exceptions to this generalist approach were limited. Only Institution B included more advanced, computer science-oriented outcomes, such as applying Bayes' rule or the minimax principle. In contrast, most institutions explicitly framed their offerings as introductory or "crash course" experiences requiring no prior technical background.

Ethical literacy was also widely represented, appearing in 9 sources and often integrated with technical instruction. Institutions framed ethical outcomes as a form of "critical literacy," encouraging learners to move beyond passive tool use. Key themes included bias awareness and mitigation (e.g., prompting for an "average American"), academic integrity (e.g., distinguishing appropriate and inappropriate uses of AI), and verification practices (e.g., identifying hallucinations or inaccuracies in AI-generated outputs).

Social learning outcomes, including understanding the historical and philosophical dimensions of AI (e.g., the Turing Test), were less frequently observed, appearing in only 6 sources. While institutional documents often referenced "societal impact" broadly, interview data suggested that these outcomes were secondary to the more immediate priorities of technical and ethical literacy.

Triangulation of Data A (documents) and Data B (interviews) identified two key areas of misalignment: Workforce Readiness and Tool Specificity. In terms of workforce readiness, a clear gap emerged between institutional messaging and instructional practice. Although 11 institutions referenced workforce preparation in their public-facing materials, only two interview participants (Institutions O and L) described learning outcomes explicitly tied to professional or industry-specific applications. For most institutions, workforce readiness functioned more as a strategic signal than as a fully operationalized curricular outcome. As one participant noted, "We talk about workforce readiness, but we're not actually designing assignments that map directly to workplace tasks."

A second area of misalignment involved institutions' emphasis on a "tool-agnostic" approach, which avoided reliance on specific platforms such as ChatGPT or Claude. Interview participants identified two primary motivations for this strategy: (a) the rapid pace of technological change, coupled with concerns about cost and equitable access; and (b) a desire to prioritize transferable AI literacy skills over platform-specific competencies. While this approach supports adaptability, it also introduces trade-offs. In particular, the absence of a standardized tool complicates the implementation of consistent accessibility checks, limiting the integration of Universal Design for Learning (UDL) principles across institutions.

Overall, findings from both Data A and Data B consistently indicate that few institutions are grounding their AI literacy learning experiences in established frameworks, and that gaps remain between institutional goals and instructional practice.

Discussion

Because AI literacy initiatives are increasingly delivered through internet-enabled environments, inconsistencies in framework adoption have direct implications for equity and the quality of these learning experiences. The findings reveal a clear disconnect between institutional messaging and instructional design. For example, while publicly available materials frequently emphasized workforce readiness as a central goal of AI literacy, this emphasis was largely absent in interview data describing actual instructional practices. This gap suggests broader challenges in articulating and translating AI literacy goals into coherent and measurable learning outcomes across institutional contexts.

At the same time, the findings indicate that, despite the availability of well-documented AI literacy frameworks, their adoption across institutions remains limited. Rather than applying standardized models, institutions tend to design AI learning experiences aligned with local contexts and priorities, often adapting existing resources or developing their own approaches. In the absence of shared frameworks or common reference points, institutions rely on informal influences, peer examples, and locally defined priorities to shape their understanding of AI literacy. As reflected in the findings, this results in a focus on immediate, pragmatic concerns—such as understanding how AI tools function, how to use them effectively, and how to address ethical risks—while broader dimensions, including social and equity considerations, receive comparatively less emphasis.

This pattern suggests that learning outcome development and framework adoption are shaped by the urgency of rapidly evolving technologies and institutional pressures to respond quickly. In such contexts, engagement with broader conceptual or equity-oriented dimensions of AI literacy may be deprioritized. Importantly, the absence of formal framework engagement is not neutral; it influences which aspects of AI literacy are foregrounded and which remain underdeveloped. Institutions may perceive existing frameworks as too rigid or insufficiently aligned with their specific contexts, or they may lack awareness of available frameworks due to limited visibility and guidance within the field.

While this localized, context-driven approach provides flexibility, it also introduces risks of inconsistency and uneven coverage. The findings highlight concerns related to framework awareness, perceived relevance, and institutional prioritization. Without a shared language or structural foundation, AI literacy initiatives vary substantially in scope and emphasis, limiting coherence across learning experiences and complicating efforts to assess outcomes. As demonstrated in the findings, such variation contributes to uneven interpretations of AI literacy and potentially uneven learner outcomes. These conditions are particularly consequential for equity-oriented and accessibility-focused practices, which require deliberate planning and coordinated design efforts to ensure meaningful and consistent integration across institutional settings.

Implications

These findings position AI literacy not only as a curricular concern, but also as a question of institutional governance, coherence, and equity in AI-enabled educational environments. They have direct implications for how institutions conceptualize, design, and govern internet-enabled AI literacy initiatives in higher education. Implications are organized across four interrelated areas related to teaching and learning.

Institutional Strategy and Policy

Calls from organizations such as EDUCAUSE emphasize the need for a shared language around AI literacy to support clearer learning outcomes, cross-institutional dialogue, and more coherent design decisions (Kassorla et al., 2024; Robert, 2025). Findings from this study indicate that, while well-documented AI literacy frameworks exist, their adoption across institutions remains limited. This reflects gaps in dissemination, perceived relevance, and institutional prioritization, and contributes to inconsistency, equity concerns, and fragmented development of AI literacy learning experiences. Institutions appear to favor context-driven, human-centered, or backward-design approaches that do not draw from standardized models, often relying on independent research, informal influences (e.g., ACM curriculum, accessibility frameworks), or peer institutions.

While this approach affords flexibility, the findings suggest that institutions may benefit from establishing shared, campus-wide AI literacy outcomes informed by existing, research-based frameworks. Such alignment can strengthen coherence while still allowing for local adaptation.

Given the documented gap between stated workforce-readiness claims (present in 11 institutions' public materials) and the limited emphasis in interviews (found only in Institutions O and L), institutions should avoid overstating workforce preparation unless it is directly reflected in enacted learning outcomes.

Online Course and Professional Development Design

Because many AI literacy initiatives are delivered through internet-enabled formats, instructional design decisions have direct consequences for coherence and learner experience. Findings indicate a reliance on one-off workshops or introductory offerings designed to respond quickly to emerging needs. While these approaches support rapid deployment, they may limit coherence across learning experiences.

Institutions may benefit from developing sequenced and scaffolded professional development pathways aligned with explicit AI literacy outcomes. Introductory or "crash course" options may remain appropriate, particularly for staff and faculty, when paired with clearly defined learning outcomes. Tiered or scaffolded pathways, including microcredentials, can further support intentional design, continuity, and alignment across learning experiences.

Equity, UDL, and Accessibility

The strong trend toward tool-agnostic design suggests that institutions should intentionally align this stance with implementation practices. Interview data indicated that tool-neutral approaches were used to address rapid technological change and mitigate cost and access barriers; however, a tool-agnostic approach also introduces challenges for consistent accessibility and Universal Design for Learning (UDL) implementation.

Institutions should ensure that the absence of a standardized tool does not result in gaps in accessibility or UDL review processes. In addition, institutions should expand policy attention to multilingual learners and LO_Social outcomes, both of which appeared inconsistently across institutions. Only one institution explicitly addressed the needs of multilingual or international learners, and LO_Social outcomes appeared more frequently in documents than in interviews, suggesting a gap between stated intentions and instructional practice.

These findings highlight the need to integrate multilingual and UDL checkpoints into course approval and review processes. Doing so can help ensure that accessibility and equity considerations are consistently addressed across instructional design. This is particularly important for supporting equitable access in internet-enabled AI literacy initiatives.

Transparency and Assessment

Limited adoption of shared frameworks complicates institutions' ability to clearly articulate and assess AI literacy goals across learning experiences. Developing a public-facing institutional AI literacy roadmap may support transparency by clarifying learning outcomes, professional development pathways, and access expectations.

Regular assessment and reporting of AI literacy outcomes may further enhance transparency, support equity monitoring, and improve coherence across internet-enabled learning initiatives.

Together, these findings suggest that strengthening coherence, equity, and intentionality in AI literacy implementation requires clearer institutional commitments, explicit alignment of learning outcomes, structured professional development pathways, deliberate integration of multilingual and accessibility considerations, and sustained attention to the gaps between stated goals and enacted instructional design.

Limitations

This study has several limitations that should be considered when interpreting the findings. First, Data A analysis relied on publicly available institutional materials, which varied widely in depth, clarity, and completeness. Because publicly posted materials differ in how fully they document design decisions, stated learning outcomes, framework use, some institutional practices may be underrepresented or missing entirely.

Interview participation was limited to nine individuals across the 16 institutions contacted, and those who volunteered may represent individuals more directly involved, more invested in the topic, or more familiar with AI literacy efforts. This creates the possibility of response bias toward institutions with more developed or visible AI initiatives.

Third, while the study employed a structured, iterative coding process, including independent coding rounds, collaborative reconciliation, and Jaccard similarity checks, qualitative coding remains an interpretive process. Despite improvements in coder agreement (from .57 to .69), discrepancies across codes indicate areas where coder interpretation varied, which may have shaped theme boundaries or emphases.

Fourth, the study's focus was limited to learning experiences designed for non-technical learners, which means the findings may not generalize to AI literacy efforts developed within computer science, data science, engineering, or other technical programs. These settings may engage more systematically with existing AI literacy frameworks or emphasize different outcome domains.

Fifth, because institutions are designing AI literacy in real time, the field is rapidly evolving. The materials and interviews captured in this study represent a moment in early adoption, and institutional practices, frameworks, and policies may change quickly as AI capabilities expand and governance structures mature.

Finally, while the study identifies gaps in multilingual support, UDL integration, social outcome development, and accessibility processes, these findings reflect what institutions made visible in documents or described in interviews, not necessarily the full scope of internal conversations, unpublished guidelines, or informal practices. As such, the absence of evidence in these domains should not be interpreted as evidence of complete institutional omission.

Conclusion

By identifying reliance on informal influences and independent research rather than formal frameworks, the findings in this study underscore systemic challenges in dissemination, perceived relevance, and institutional prioritization. While assembling elements they deemed necessary for AI literacy affords institutional flexibility, it also raises concerns about consistency, equity, and the long-term coherence of AI literacy development across higher education. Among the institutions studied, they emphasized technical and ethical competencies, giving less attention to the social dimensions of AI. This imbalance suggests that current AI literacy efforts tend to prioritize functional use and ethical caution over broader social, civic, or global considerations. Without deliberate alignment between guiding frameworks and articulated learning outcomes, institutions risk fragmented approaches to AI literacy that hinder coherence, limit meaningful assessment, and obscure what students are ultimately intended to understand about AI's role in society.

There is limited value in institutions reinventing AI literacy in isolation. Rather than relying on ad hoc or reactive design decisions, institutions would benefit from examining existing, empirically grounded AI literacy frameworks and adapting those that align with their institutional contexts, values, and learner populations. Learning outcomes should be intentionally derived from these frameworks to clearly inform instructional design, implementation, and assessment.

These findings highlight that AI literacy development in higher education is not solely a question of curricular design or instructional strategy, but reflects broader institutional processes of governance, prioritization, and sense-making under conditions of rapid technological change. Across the examined institutions, decisions about AI literacy learning outcomes, frameworks, and delivery formats were shaped by local contexts, perceived urgency, and available expertise rather than by shared sector-wide guidance. This decentralized approach enabled institutions to respond quickly to emerging AI capabilities, yet it also produced uneven attention to accessibility, equity, and coherence across learning experiences.

The analysis further suggests that AI literacy initiatives function as socio-technical systems, where choices about tools, frameworks, and pedagogical emphasis are intertwined with organizational structures and institutional values. Emphases on functional use and ethical caution, while important, often overshadowed broader social or civic dimensions of AI, raising questions about what kinds of AI literacies are being cultivated and for whom. In this way, AI literacy becomes a site where institutional goals related to workforce readiness, risk management, and educational inclusion converge and, at times, conflict.

By examining AI literacy implementation at the institutional level, this study contributes to understanding how educational systems are navigating AI adoption beyond individual classrooms. The findings underscore the importance of intentional governance structures and clearer articulation of learning outcomes to support more equitable, coherent, and transparent AI literacy efforts across higher education.

Declaration of Generative AI and AI-Assisted Technologies in the Manuscript Preparation Process

During the preparation of this manuscript, the authors used Microsoft Copilot as an assistive analytic tool to support qualitative analysis tasks, including early codebook organization, identification of potential coding discrepancies, and calculation of semantic Jaccard similarity scores. Copilot was used to generate preliminary code structures and flag areas of divergence across coded data; however, all outputs were reviewed, verified, and interpreted by the research team.

Analytic decisions including code definitions, revisions, theme development, and resolution of discrepancies were made exclusively by the authors through collaborative discussion and iterative review. Copilot functioned as a support tool rather than an analytic decision-maker. The authors take full responsibility for the content of the published article.

References

Akhmadieva, R. S., Lysenko, V. V., Guryanova, I. V., Nemtsev, A. N., & Khairullina, E. R. (2025). Research trends in the use of artificial intelligence in higher education: A bibliometric analysis. *Frontiers in Education,* 9, Article 1438715. https://doi.org/10.3389/feduc.2024.1438715

Almatrafi, O., Johri, A., & Lee, H. (2024). A systematic review of AI literacy conceptualization, constructs, and implementation and assessment efforts (2019–2023). *Computers and Education Open*, *6*, 100173. https://doi.org/10.1016/j.caeo.2024.100173

Chiu, T. K. F., Choi, A., & Lin, T.-J. (2024). What are artificial intelligence literacy and competency? A comprehensive framework to support them. *Computers and Education, Open*, 6, 100171. https://doi.org/10.1016/j.caeo.2024.100171

Daher, R. (2025). Integrating AI literacy into teacher education: A critical perspective paper. *Discover Artificial Intelligence*, *5*(1), 217. https://doi.org/10.1007/s44163-025-00475-7

Deterding, N. M., & Waters, M. C. (2021). Flexible coding of in-depth interviews: A twenty-first-century approach. *Sociological Methods & Research*, *50*(2), 708–739. https://doi.org/10.1177/0049124118799377

Hibbert, M., Altman, El., Shippen, T., & Wright, M. (2024). *A framework for AI literacy*. EDUCAUSE Review. https://er.educause.edu/articles/2024/6/a-framework-for-ai-literacy

Horowitz, H., & Foulkes, J. (2025). The ICAIL framework: A critical and inclusive approach to AI literacy for non-technical learners in higher education. *Innovations in Learning*, *1*(1), https://learninginnovations.lib.buffalo.edu/articles/11

Jin, Y., Martinez-Maldonado, R., Gašević, D., & Yan, L. (2024). *GLAT: The Generative AI Literacy Assessment Test* (arXiv:2411.00283). arXiv. https://doi.org/10.48550/arXiv.2411.00283

Kassorla, M., Georgieva, M., & Papini, A. (2024, October 17). *AI literacy in teaching and learning: Executive summary*. EDUCAUSE. https://www.educause.edu/content/2024/ai-literacy-in-teaching-and-learning/executive-summary#ExecutiveSummary

Kennedy, K., & Gupta, A. (n.d.). *AI & Data Competencies: Scaffolding holistic AI literacy in Higher Education* (arXiv:25010.24783). arXiv. https://doi.org/10.48550/arXiv.2510.24783

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–16. https://doi.org/10.1145/3313831.3376727

Mills, K., Ruiz, P., Lee, K., Coenraad, M., Fusco, J., Roschelle, J., & Weisgrau, J. (2024). *AI literacy: A framework to understand, evaluate, and use emerging technology*. Digital Promise. https://doi.org/10.51388/20.500.12265/218

Ng, D. T. K., Leung, J. K. L., Chu, S. K. W., & Qiao, M. (2021). Conceptualizing AI literacy: An exploratory review. *Computers and Education: Artificial Intelligence*, 2, 100041. https://doi.org/10.1016/j.caeai.2021.100041

Rismanchian, S., & Doroudi, S. (2025). The evolution of research on AI and education across four decades: Insights from the AI×Ed framework. *International Journal of Artificial Intelligence in Education*. *35*, 2797-2820. https://doi.org/10.1007/s40593-025-00483-9

Robert, J. (2025). *2025 horizon action plan: Building skills and literacy for teaching with GenAI*. EDUCAUSE. https://library.educause.edu/resources/2025/9/2025-educause-horizon-action-plan-building-skills-and-literacy-for-teaching-with-genai

Southworth, J., Migliaccio, K., Glover, J., Glover, J., Reed, D., McCarty, C., Brendemuhl, J., & Thomas, A. (2023). Developing a model for AI Across the curriculum: Transforming the higher education landscape via innovation in AI literacy. *Computers and Education: Artificial Intelligence*, *4*, 100127. https://doi.org/10.1016/j.caeai.2023.100127

Touretzky, D., Gardner-McCune, C., & Seehorn, D. (2022). Machine learning and the Five Big Ideas in AI. *International Journal of Artificial Intelligence in Education, 33*, 223-266. https://doi.org/10.1007/s40593-022-00314-1

UNESCO. (2024). *AI competency framework for students*. https://doi.org/10.54675/JKJB9835

Valentini, A., & Blancas, A. (2025). *The challenges of AI in higher education and institutional responses: Is there room for competency frameworks*? UNESCO International Institute for Higher Education in Latin American and the Caribbean. https://www.iesalc.unesco.org/en/articles/challenges-ai-higher-education-and-imperative-competency-frameworks