

# Guidelines for Designing AI Technologies to Support Adult Learning

Jennifer M. Reddig\*  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Wesley G. Morris  
Vanderbilt University  
Nashville, Tennessee, USA

John Kos  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Momin Naushad Siddiqui  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Sandeep Kakar  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Min Kyu Kim  
Georgia State University  
Atlanta, Georgia, USA

Glen R. Smith Jr.\*  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Yoojin Bae  
Georgia State University  
Atlanta, Georgia, USA

Rahul K. Dass  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Daniel Weitekamp  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Alex Endert  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Chris Dede  
Harvard University  
Cambridge, Massachusetts, USA

Christopher J. MacLellan  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Sanaz Ahmadzadeh Siyahrood  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Kaitlyn Crutcher  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Jinho Kim  
Georgia State University  
Atlanta, Georgia, USA

Ploy Thajchayapong  
Georgia Institute of Technology  
Atlanta, Georgia, USA

Scott Crossley  
Vanderbilt University  
Nashville, Tennessee, USA

Ashok Goel  
Georgia Institute of Technology  
Atlanta, Georgia, USA

## Abstract

AI-powered educational technologies have demonstrated measurable benefits for learners, but their design and evaluation have largely centered on K-12 contexts. As a result, many AI-supported learning systems remain poorly aligned with the needs, constraints, and goals of adult learners. To better understand how AI systems function in adult education, this paper examines the deployment of several AI learning technologies developed within a multidisciplinary, national research institute in the United States focused on adult learning and online education. Drawing on longitudinal deployment data, we conducted a reflexive thematic analysis to identify recurring challenges and design considerations across systems. These insights were synthesized into a set of 19 design guidelines intended to inform future AI-supported adult learning technologies. We demonstrate the utility of these guidelines through a heuristic

evaluation of the deployed systems. Lastly, we present a guideline exploration tool that aids in the ideation of technologies by connecting the guidelines to stakeholder statements surfaced in the analysis process.

## CCS Concepts

• **Human-centered computing** → **Heuristic evaluations.**

## Keywords

Artificial Intelligence, Design Guidelines, Adult Learning, Andragogy, Educational Technology

## ACM Reference Format:

Jennifer M. Reddig, Glen R. Smith Jr., Sanaz Ahmadzadeh Siyahrood, Wesley G. Morris, Yoojin Bae, Kaitlyn Crutcher, John Kos, Rahul K. Dass, Jinho Kim, Momin Naushad Siddiqui, Daniel Weitekamp, Ploy Thajchayapong, Sandeep Kakar, Alex Endert, Scott Crossley, Min Kyu Kim, Chris Dede, Ashok Goel, and Christopher J. MacLellan. 2026. Guidelines for Designing AI Technologies to Support Adult Learning. In *Designing Interactive Systems Conference (DIS '26)*, June 13–17, 2026, Singapore, Singapore. ACM, New York, NY, USA, 22 pages. <https://doi.org/10.1145/3800645.3813102>

\*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License. *DIS '26, Singapore, Singapore*

© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2563-0/2026/06  
<https://doi.org/10.1145/3800645.3813102>

## 1 Introduction

AI-powered educational technologies enable adaptive, data-driven instruction across a wide range of learning contexts and have demonstrated measurable benefits for learner engagement, personalization, and instructional efficiency at scale [14]. Although recent trends in educational technology research reflect a growing emphasis on higher education and adult learning populations [6, 43], many AI-supported systems continue to be designed and evaluated within K-12 settings, where learning typically occurs in formal classrooms with relatively stable schedules, curricula, and support structures. A systematic review by Bernacki et al. [5] suggests that research on personalized learning has largely overlooked adult learning and workplace training populations. As a result, prevailing design knowledge, assumptions, and system evaluation criteria are tightly coupled to the needs and constraints of younger learners and traditional schooling contexts [23]. This renders many AI-supported learning technologies poorly aligned with the needs of adult learners, who increasingly pursue education through professional training and informal learning pathways [56].

Adult learners differ from younger learners not only demographically, but motivationally and contextually, with learning situated within the broader constraints of their daily lives. Adult learning is often self-directed and goal-oriented, shaped by immediate needs such as career advancement, credentialing, or reskilling [44, 45]. At the same time, adult learners must balance education alongside employment, family responsibilities, and other social obligations, which constrains when, where, and how learning occurs. These conditions shape patterns of technology use and challenge design assumptions that presuppose continuous participation, uniform goals, or stable instructional environments. Consequently, learning technologies intended for adult learners must account for the contextual conditions under which learning takes place, rather than inheriting design patterns optimized for K-12 settings.

This paper examines these challenges through an empirical investigation of AI-powered learning technologies deployed in adult education contexts. Our analysis focuses on a set of systems developed and deployed within the National AI Institute for Adult Learning and Online Education (AI-ALOE) in the United States [30]. Even though these technologies were implemented and deployed across diverse adult learning environments, stakeholders like learners, instructors, administrators, and technology development teams repeatedly surfaced similar needs, challenges, and design considerations. The presence of these patterns across otherwise distinct deployments motivated a more systematic, cross-technology synthesis of findings.

Drawing on longitudinal deployment data—including focus groups with instructors and learners, technical artifacts, and progress reports—we analyze how these technologies were designed, implemented, and experienced in practice. Using reflexive thematic analysis [9], we synthesize themes and insights from these deployments into a set of 19 design guidelines. We then demonstrate the applicability of these guidelines by evaluating them against the same set of technologies, highlighting opportunities for improvement and the trade-offs that emerge in real-world adult learning environments. Collectively, this work contributes:

- (1) An empirically grounded set of 19 guidelines for AI-powered adult learning technologies,
- (2) A method for conducting heuristic evaluations of adult learning technologies using these guidelines, and
- (3) A guideline exploration tool that informs meaningful ideation and technology refinement by linking guidelines to stakeholder statements surfaced during our analysis.

We begin by situating this work within prior research on adult education and AI-supported learning technologies. We then describe the learning technologies examined in this study and the adult learning contexts in which they were deployed. Next, we detail our research methodology, including data sources and analytic procedures. We subsequently present the resulting design guidelines, followed by a discussion of their implications for the design and evaluation of learning technologies for adult learners and directions for future research.

## 2 Related Work

Knowles [44] introduces andragogy, the theory of adult learning, which points out that adults are independent, self-directed learners with years of life experience who require meaningful, problem-oriented teaching strategies. In contrast, Knowles describes pedagogy, the theory of child learning, as being characterized by a more dependent relationship between instructor and child, where the instructor largely shapes the direction and content of the curriculum. According to Knowles, adults require instructional strategies that are collaborative and demonstrate how content supports their learning goals. Adults exist in a different social and cultural space than children, distinguished by a greater degree of self-directed decision-making. Educational technology for adults must support a different set of needs than tools designed for children [79]. Though recent efforts have explored AI tools in higher education [52], these technologies do not take into account the unique social contexts and outside demands placed on adult learners.

While andragogy characterizes adult learners, pedagogy offers many well-established instructional strategies and cognitive principles that are also applicable to adult learning, such as scaffolding [84], formative feedback [72], guided reflection and experimentation [60], and structured learning activities [7]. These practices are general across both K12 and adult learner contexts. When applied to adult learning, pedagogical strategies should be adapted to respect learner autonomy, leverage prior experience, and align with goal-oriented motivations [44, 79]. In this way, pedagogy is a foundational set of instructional design principles that can be translated to adult learning contexts.

Particularly relevant to adult online learning, Garrison et al. [27] introduced the Community of Inquiry (CoI) framework to investigate online learning environments, which suggests that a meaningful educational experience occurs through the interaction of teaching presence, social presence, and cognitive presence. Cognitive presence is the learner's internal engagement with the learning experience. Social presence is the group collaboration and discussion toward learning outcomes, and allowing learners to project their authentic selves. Teaching presence is the organizational structure, design, and facilitation of the learning environment. Social

presence is particularly critical in online contexts, where asynchronous and distributed interaction can inhibit the development of interpersonal relationships. Unlike face-to-face settings where social connections form naturally, online learning environments require deliberate design to foster social presence, as its absence has been linked to learner isolation and disengagement [29]. The three presences are interconnected: teaching presence creates the conditions for social and cognitive presence to develop and interact, cognitive presence is deepened by collaborative inquiry, and social presence builds motivation and trust for sustained effort. Adult and online learning environments require intentional design and support for these elements, since adult learners often engage in distributed, asynchronous, and professional learning contexts [28, 56]. The technologies developed by AI-ALOE are designed to support different elements and combinations of elements within the CoI framework. By looking at the seven tools developed within the institute collectively, we offer a view into how meaningful learning can be supported by an ecosystem of online educational technologies.

Adult learning in online environments involves challenges related to sustained engagement, feedback, and facilitation interaction across distributed and asynchronous contexts. Artificial intelligence has increasingly been explored as a way to support teaching and learning [12, 14]. Prior research shows that AI-supported educational tools can assist with individualized feedback, content adaptation, and instructional coordination across pedagogical theories [22]. However, empirical studies have also identified several challenges with AI-supported learning, including learner frustration and tools that are difficult to understand or use effectively [15, 85]. In response, several frameworks have been proposed to guide the use of AI in education, including pedagogically oriented classifications [22] and principle-based approaches focused on transparency, agency, and accountability [24]. Despite their contributions, these frameworks often remain high-level and offer limited guidance for translating pedagogical principles into the design of AI-supported tools, particularly in adult learning contexts. Efforts to adapt broader educational design frameworks, such as Universal Design for Learning, to AI-supported contexts have faced similar challenges related to complexity and limited empirical support [88].

The tools developed within AI-ALOE span several established categories of educational technology. Intelligent tutoring systems have been widely studied, with learning gains approaching those of one-on-one human tutoring [80]. Interactive video platforms embed questions and tasks into educational videos, which improves learning outcomes over passive viewing [71]. AI-powered chatbots have grown substantially by applying large language models to provide on-demand instructional support and personalized feedback, though there are ongoing concerns around reliability and academic integrity [48]. Intelligent textbooks also have integrated large language models to generate formative feedback, assessment questions, and interactive dialog grounded in textbook content [76]. Scientific modeling tools allow students to construct, test, and revise hypotheses through interactive simulations [19, 31]. Social agents are a more recent innovation, designed to combat student isolation in online learning by suggesting potential collaborators with shared interests and goals [81]. Since the suite of AI-ALOE tools spans a broad range of techniques, learning goals, and strategies, examining the constraints and impacts of the tools collectively

provides a more well-rounded perspective of AI-supported learning tools.

Within human–AI interaction research, design frameworks and guidelines have played an important role in translating high-level principles into actionable support for system design. For example, the human-AI interaction guidelines proposed by [1] demonstrates how empirical evidence can be consolidated into reusable guidelines that informs evaluation and design practice. More recent research has emphasized the value of design artifacts like toolkits, cards, and structured resources to support ideation and collaboration when designing AI systems [87]. A growing body of work in AI and education has proposed guidelines and frameworks to support the responsible and effective use of AI in learning and assessment [65, 66, 68]. Recent synthesis shows that much of this guidance is framed at the policy level, rather than system guidance [26]. This literature highlights the importance of empirically grounded design guidance that connects pedagogical theory with the practical realities of designing, evaluating, and iterating on AI-powered educational systems for adult learners.

## 3 Background

### 3.1 AI-ALOE

The National AI Institute for Adult Learning for Online Education (AI-ALOE) is a large, interdisciplinary research institute formed through a diverse coalition of research-intensive and regional universities, and community and technical college systems. The institute brings together researchers from artificial intelligence, learning sciences, cognitive science, education, linguistics, and human-computer interaction around a shared mission of improving the quality and accessibility of online and hybrid education for adult learners. AI-ALOE focuses specifically on adults engaged in re-skilling, up-skilling, or continuing education [30]—learners whose educational experiences are shaped by limited time, varying prior knowledge, and the need to integrate learning into complex personal and professional lives. In pursuit of this mission, the institute develops and studies AI-driven learning technologies intended to support flexible, self-directed learning pathways, while accounting for the social, cognitive, and institutional constraints that characterize adult education contexts.

We selected AI-ALOE technologies for investigation because they offer a combination of contextual diversity and functional breadth. First, the institute’s technologies span a wide range of adult learning contexts and populations, including rural and urban settings, community colleges and research universities, and disciplines across STEM and the humanities. Second, the portfolio of tools reflects broad coverage of instructional roles and AI capabilities, ranging from tutoring and feedback to social support, metacognitive scaffolding, conceptual modeling, and instructor-facing analytics. This diversity enables analysis across multiple forms of adult learning support. Lastly, these technologies have been examined through sustained, longitudinal studies in authentic adult learning environments, generating rich data from learners, instructors, and design teams over time. Together, these characteristics make AI-ALOE an especially well-suited setting for synthesizing design guidelines that are grounded in practice and sensitive to the realities of adult learning rather than isolated or short-term laboratory studies.

## 3.2 AI Technologies

In this section, we describe each AI technology we examined as part of this study. The technologies were developed by research teams from several academic institutions, including Georgia Institute of Technology, Georgia State University, Harvard University, University of North Carolina at Greensboro, and Vanderbilt University.

**3.2.1 Apprentice Tutors.** A web-based, intelligent tutoring system (ITS) platform designed to support adult learners in practicing domain-specific procedural skills through interactive problem solving [34]. Tutors can be created in one of two ways: (1) manually, by conducting a cognitive task analysis [50] with instructors to inform the design of each tutor, or (2) through a tutor builder component that enables educators to construct interfaces via drag-and-drop interactions and generate expert models with the support of an AI agent [75]. As learners engage with dynamically generated problems, the system provides real-time correctness feedback, allowing them to iteratively refine their understanding and skills. Instructor- and learner-facing dashboards further support the learning process by visualizing skill mastery and progress over time, helping users monitor performance and make informed instructional or self-regulatory decisions.

**3.2.2 iTELL.** Intelligent Texts for Enhanced Lifelong Learning (iTELL) is an AI-powered platform that helps adult learners practice active and metacognitive reading strategies by paraphrasing and reflecting on key ideas from the material [61]. iTELL ingests instructional text in PDF format and transforms it into an interactive web application that includes content such as short constructed-response questions, page-level summaries, and end-of-page cloze test activities. iTELL's AI module provides formative feedback for learners [62, 63] and its analytics dashboard provides actionable insights for educators to aid in instructional decision-making [16, 17].

**3.2.3 Ivy.** An interactive video-based AI coaching system embedded in online courses to support procedural skill learning [18, 51]. Ivy generates structured explanations about “how” and “why” procedural steps work by integrating symbolic representations of goals and problem decompositions, and states and causal transitions, with LLM-based synthesis of explanations. These representations act as active scaffolds that shape both what is explained and how explanations are organized, supporting adult learners' understanding of complex, multi-step procedures.

**3.2.4 Jill Watson.** A retrieval-augmented generation powered generative-AI-based teaching assistant and cognitive partner that converses with learners using instructor-approved course materials [32, 39]. It aims to help adult learners access instructional support anytime, anywhere by increasing teacher presence [49] and to deepen cognitive engagement with course content by offering a conversational interface for clarification and reinforcement [53]. It supports both retrieval of text and images in response to content-based and logistical queries in both text and audio formats [78]. The system has been deployed across multiple courses in different subjects and has been integrated with existing learning management systems [39].

**3.2.5 SAMI.** A social agent embedded within online discussion platforms to encourage class participation and social interaction

[82]. It reads students' posts on the discussion forum, uses LLMs to build knowledge graphs based on student attributes, and uses graph matching to generate recommendations for social connections [38]. It also provides explanations of its reasoning and recommendations. The tool is designed to increase a learner's sense of social belonging and is used to support emotional well-being and self-efficacy. It aims to help adult learners feel more connected to their peers and instructors.

**3.2.6 SMART.** Student Mental Model Analyzer for Research and Teaching (SMART) provides AI-powered formative feedback on learner-generated summaries and supports iterative revision at a learners' own pace [20]. It integrates knowledge-based AI with generative AI to deliver multi-modal feedback, including a visualized knowledge map that highlights key concepts learners have identified or omitted [36, 42]. Through this AI-augmented summarization process, adult learners are better supported in comprehending core concepts of the material [4, 41].

**3.2.7 VERA.** A conceptual modeling tool that guides learners through the process of building, simulating, and revising scientific models [2, 3]. It provides a visual modeling interface connected to a backend simulation engine and is supported by AI-generated coaching and feedback [10]. The tool is used in courses that emphasize systems thinking and inquiry-based learning. It aims to support adult learners by helping them understand and manipulate complex systems through hands-on model construction and experimentation, scaffolding them through any missing domain knowledge they may have [46, 47].


## 3.3 AI Technology Deployments

The AI-ALOE technologies were deployed across a diverse set of adult learning environments, spanning credit-bearing and non-credit courses offered in fully online, hybrid, and in-person formats. These deployments took place across multiple academic institutions and disciplinary contexts, including nursing, algebra, computer science, ecology, web development, and research methods. Table 1 summarizes the subjects, learner age ranges, and geographic contexts of each technology's deployments, and Figure 1 illustrates how institution types connect to technologies and delivery modalities.

Several deployments focused on supporting social engagement and community building in online courses, where adult learners may experience isolation or reduced opportunities for peer interaction. For example, SAMI was embedded within online discussion platforms in self-paced courses to encourage participation, foster social connection, and support learners' sense of belonging [38]. In contrast, tools such as VERA and IVY were deployed in classroom-based and hybrid courses to support conceptual understanding and procedural learning through modeling, simulation, and structured explanation [18? ]. Apprentice Tutors was integrated into mathematics and nursing courses to provide targeted, on-demand practice problems with immediate feedback, enabling learners to develop problem mastery outside of limited class time [35]. iTELL was used in web development and research methods courses to support active reading, comprehension, and metacognitive strategies as learners engaged with complex instructional texts [17]. Across deployments, instructional goals included strengthening teaching

Technology	Subjects	Age Range	Geographic Context
Apprentice	Algebra, Nursing	21–65	Online
iTELL	Employment Regulations, Intro to Programming, Learning Analytics, Mobile App Dev, NLP, Pharmacology, Research Methods in Psychology	18–50	Online, Urban
Ivy	Knowledge-based AI	21–65	Online
Jill Watson	Accounting, Computer Science, English, Information Technology, Leadership & Management, Medicine for Business, Mobile App Dev, Nursing	20–35	Online, Rural, Urban
SAMI	Computer Science	20–50	Online, Urban
SMART	Biology, Computer Science, Education, English, Nursing	~20s	Online, Urban/Suburban
VERA	Biology, Cognitive Science, Computer Science, Ecology, Wildlife Management	18–65	Online, Rural, Urban

**Table 1: Deployment contexts of AI-ALOE technologies, summarizing the subjects, learner age ranges, and geographic settings in which each technology was used.**



**Figure 1: Connections between institution types (research universities and technical colleges), AI-ALOE technologies, and delivery modalities (online, in-person, hybrid).**

presence, increasing learner engagement, improving feedback quality, and supporting self-regulated learning.

A defining characteristic of these deployments was the close collaboration between research teams, instructors, and learners throughout design and refinement. Many teams adopted participatory and iterative approaches, working directly with faculty to align system behavior with instructional goals. For instance, the Apprentice Tutors team met regularly with instructors to conduct cognitive task analyses and identify domain-specific tutoring needs,

resulting in tutors that were tightly aligned with course content and instructional intent [35]. The SMART team introduced automated feedback tools to instructors who were new to AI-assisted instruction, providing structured onboarding and continuing support as instructors gradually integrated AI-generated feedback into their courses [41]. These collaborations allowed instructors to shape how AI features were introduced and framed for learners, while enabling research teams to adapt system designs in response to instructional constraints.

Across multiple deployments, the technologies demonstrated measurable improvements in learner outcomes and stakeholder adoption. Use of Jill Watson was associated with higher average course grades and improved retention compared to control sections [39]. Deployments of SMART showed that learners receiving AI-supported concept feedback produced higher-quality written work [4]. Apprentice showed higher in-class assessment performance after using the tutor [33]. iTELL demonstrated improvements in summary writing quality and sustained engagement across diverse adult learner populations in both hybrid and online contexts [16]. Lastly, results from VERA suggested gains in the complexity of learners' mental models, though outcomes varied by course structure and instructional emphasis [47].

Viewed collectively, the deployments surfaced a set of recurring themes that appeared across technologies, disciplines, and instructional formats. Instructors and learners consistently highlighted similar challenges and points of friction during use, regardless of the specific system involved. The presence of these patterns across otherwise distinct deployments suggested that they reflected broader features of adult learning experiences. These observations motivated the cross-technology analysis presented in this paper, forming the empirical foundation for the design guidelines articulated in the following sections.

## 4 Methodology

To conduct our synthesis, we performed a *reflexive thematic analysis* [8, 9] of user and stakeholder feedback. The goal was to support future development of AI-powered education tools by highlighting common design considerations, user needs, and contextual factors that emerged across our multiple technologies and deployments. This section details how we compiled the relevant data and performed the analysis.

### 4.1 Data Source Collection

We began our analysis by assembling a corpus of qualitative data that captured how the AI-ALOE technologies were designed, deployed, and experienced in adult learning contexts. These data were drawn from multiple sources to reflect both the intentions of research teams and the usage experiences of learners and instructors. Specifically, we collected a total of 15 presentations from the research and development teams of the technologies that documented design rationales, relevant learning science frameworks, and instructional considerations; transcripts from 17 focus groups conducted with learners, instructors, and administrators; and 3 sets of cross-team feedback from NSF program committee meetings.

To support a systematic, cross-technology analysis, we transformed these heterogeneous materials into a standardized data format. From presentations, reports, and design documents, we extracted discrete statements reflecting educational concepts such as learning science theories and design considerations articulated by research teams. From focus group transcripts, we extracted individual thoughts, experiences, and reflections expressed by learners, instructors, and administrators, preserving the original phrasing to maintain context. Each extracted item was entered as a single row in a digital spreadsheet and annotated with metadata including the associated technology, the raw text excerpt, and the source context

(e.g., “SMART - Instructor Focus Group”, “iTELL - Design Presentation”, etc.). This standardized representation enabled us to assign unique identifiers to each data item (e.g., “SMART-Learner Focus Group-299”), allowing insights to be traced back to their original source throughout analysis if additional context was required. This structured dataset formed the foundation for the reflexive thematic analysis described in the following section.

### 4.2 Reflexive Thematic Analysis

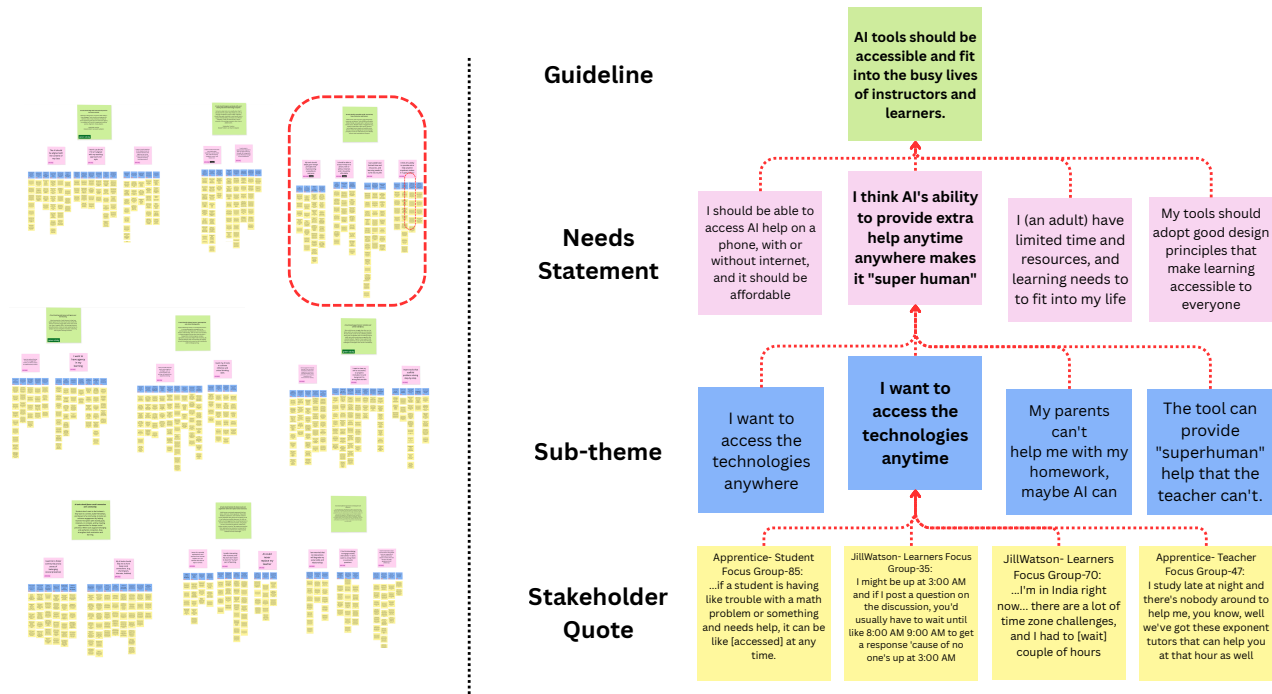
We conducted a reflexive thematic analysis [8, 9] to synthesize patterns across the collected data. The analysis was carried out by a deliberately diverse research team that included at least one representative from each of the technology design teams, one researcher who led and conducted the focus groups, and 3 additional HCI researchers who were not part of the development or focus group teams. To ensure a comprehensive overview, all team members participated in a collaborative synthesis process to analyze all of the data following the procedure of Holtzblatt [37]. Team members represented over five institutions and brought complementary perspectives from computer and learning sciences, educational technology, and human-computer interaction. This composition was intended to ensure a plurality of perspectives informed the analysis, spanning design intent, deployment experience, and learner- and instructor-facing observations. Through iterative discussion and debate, the team collaboratively identified, refined, and synthesized themes, drawing on both the data and their situated expertise with adult learning technologies. As a result, we acknowledge that it is possible another research team, drawing on different experiences and values, might arrive at an alternate set of themes or guidelines when analyzing the same dataset. Indeed, we selected this approach because the goal of this study was not to establish a single objective account of stakeholder experiences, but to surface recurring design considerations grounded in the situated experiences of learners, instructors, and technology developers across multiple AI systems and contexts.

All extracted data items—approximately 1,600 individual statements—were imported into Miro<sup>1</sup> and initially organized by technology. Each item was represented as a yellow sticky note containing a verbatim quote, observation, or design insight from the dataset. Following an affinity diagramming process, the team collaboratively grouped related yellow notes into first-level thematic clusters, each labeled with a blue sticky note (See Figure 2). These blue notes captured shared patterns across small sets of related items (typically four to six) and were framed in the first-person (“I”) perspective to reflect the dominant stakeholder voice represented in the underlying data (e.g., learner, instructor, or developer). For example, a blue note such as “*I want students to develop transferable critical thinking skills*” reflected a recurring instructor-oriented concern present across multiple data sources.

Next, the team grouped related blue notes into higher-level pink notes, also written in the “I” perspective, but reframing broader concerns of multiple stakeholders as desired features. For instance, a pink theme such as “*I want to design AI tools that are intuitive, easy to use, and promote natural interactions to avoid confusion over the interface and expectations of the tool*” encompassed learner

<sup>1</sup><https://miro.com/>

## Affinity Diagramming Process



**Figure 2: Affinity diagramming process used to derive design guidelines from qualitative data. Stakeholder quotes (yellow notes) were clustered into sub-themes (blue notes) capturing recurring perspectives. Sub-themes were synthesized into needs statements (pink notes) spanning multiple stakeholder groups. Finally, needs statements informed distilled design guidelines (green notes).**

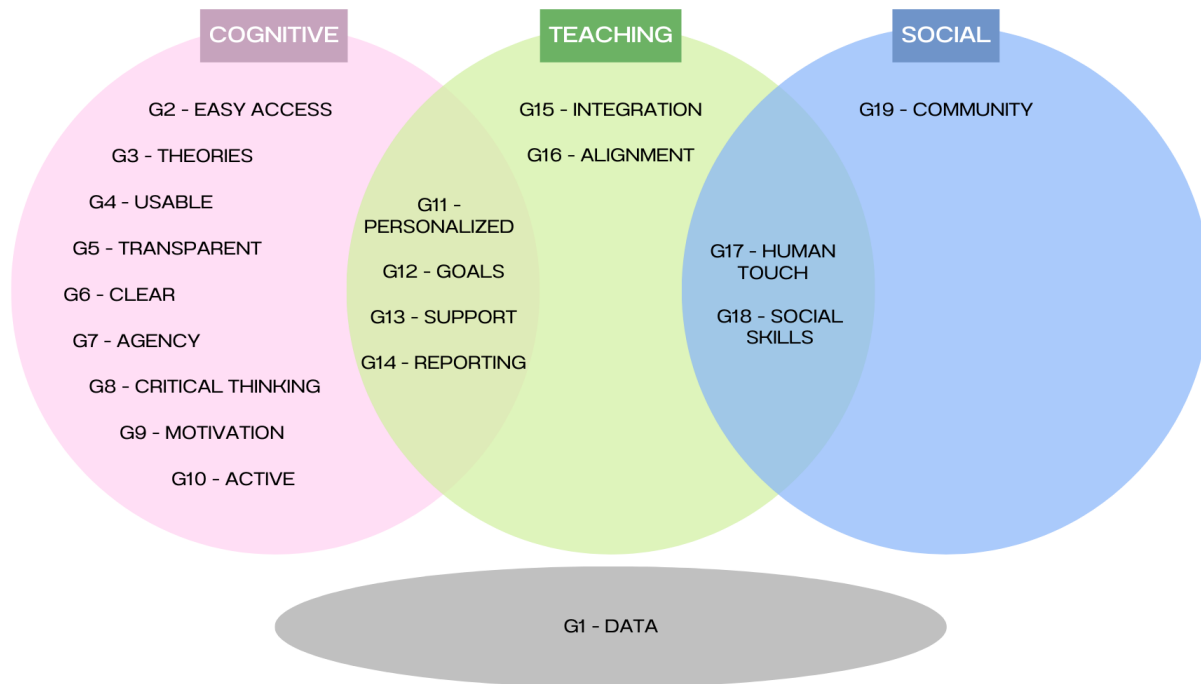
anxieties about data deletion ("I'm afraid of accidentally deleting my work"), instructor concerns about usability ("I make mistakes because of the design of the UI"), and developer reflections on overall interface design. Throughout this process, themes were iteratively discussed, merged, split, or reframed as the team revisited the data and reflected on alternative interpretations.

Finally, the pink themes were synthesized into a set of design guidelines, represented as green sticky notes and framed as actionable statements beginning with "AI tools should...". This synthesis involved a collaborative review where the entire research team examined each thematic group of pink notes, proposed candidate guidelines, and iteratively revised the wording until reaching a consensus that each guideline accurately reflected the consolidated stakeholder voice. Each guideline was accompanied by a detailed description and further annotated along several dimensions, including the primary stakeholders affected, the types of technology problems addressed (e.g., access, trust, engagement), and the Community of Inquiry dimensions—cognitive, teaching, and social presence [27]—that the guideline supports. These annotations supported later analysis and application of the guidelines, including the heuristic evaluation described in Section 6.

## 5 Design Guidelines

Table 2 shows the 19 guidelines that resulted from our analysis. Each guideline represents a systematic distillation of the multi-level thematic groupings described in our methodology, moving from specific stakeholder concerns to higher-order design principles. By grounding each guideline in the recurring needs and friction points surfaced across various adult learning contexts, we ensure they remain actionable and relevant to the practical challenges of educational technology design. As described earlier, AI-ALOE adopted Garrison et al. [27]'s Community of Inquiry as a core framework to guide the design of its learning tools. Accordingly, these guidelines are labeled according to the associated Community of Inquiry element. Each guideline is also labeled by the stakeholder most directly affected if it is not satisfied, as well as by the potential technological issues that could be mitigated through its application. Figure 3 illustrates how the guidelines overlap with Community of Inquiry dimensions. In the supplementary materials, we provide the full affinity diagram, including the stakeholder needs statements and supporting themes behind each guideline.

We only labeled the stakeholders who would experience an immediate impact if the guideline is not addressed. For example, guideline G4 says that "AI tools should be easy to understand and frictionless to use". Making a tool easy to understand will be an immediate benefit to the end users, the learners and instructors. If the tool



**Figure 3: Design guidelines organized across three overlapping Community of Inquiry dimensions: Cognitive (learner experience and understanding), Teaching (instructional design and alignment), and Social (community and interpersonal development). Shared regions indicate principles that bridge domains, while G1 (Data Practices) sits outside the CoI framework.**

is not easy to understand, users may stop engaging with the tool. Though the researcher is not immediately impacted by this issue, they will be affected in the future when they are unable to collect research data due to low engagement and adoption. We chose not to include the researcher in the stakeholders because, with downstream effects, nearly every stakeholder could be impacted by every guideline which would lead to a trivial taxonomy. To preserve the usefulness and clarity of the guidelines, we labeled each guideline with only the immediately impacted stakeholders.

In our supplementary materials, we also present the low-level themes and mid-level themes that support each design guideline in the full affinity diagram. We believe that by reviewing the themes surfaced from concerns of stakeholders, future researchers and developers can generate practical solutions to meet the guidelines based on specific real-world concerns. In addition, reviewing the supporting themes could help apply the guidelines to broader contexts, beyond the ones from which they were derived.

Across the 19 guidelines, we observed clear differences in how concerns were distributed among stakeholder groups (Figure 4). Instructors most frequently commented on the usability of AI tools (G4) and students most frequently commented on the amount of educational support provided (G13), but both stakeholders commented heavily on these two dimensions. Instructors also frequently highlighted the alignment of AI tools with their personal instructional approach (G16) because using AI tools with contrasting approaches

can be a challenge. Student comments were also particularly concentrated on guidelines addressing building community (G19) and social competencies (G18), indicating a strong concern for how AI will shape peer and professor interaction. In contrast, researcher comments were more prominent in guidelines tied to instructional best practices (G8, G9, G10) and using data for learning engineering (G14). The category of Learning Theories (G3) was also only a concern for researchers and not mentioned by the other stakeholders. While several guidelines received equal attention from all three groups (G2, G4, G5, G13), the uneven distribution demonstrates that different stakeholders prioritize different dimensions of AI-supported learning. The guidelines collectively provide a broad representation of the concerns raised across all stakeholder groups.

Guideline **G1 (Data Practices)** was not assigned a Community of Inquiry component. This guideline was derived mainly from student concerns about not knowing what data is being collected, how it is stored, and who can access it. Since aligning AI technologies with G1 would not increase student cognitive, social, or teaching presence, we elected not to force G1 into the Community of Inquiry framework. Students expressed a desire to have meaningful control over their data, to review, manage, or delete it. AI technologies that use third-party Large-Language Models need to be transparent with users in an educational setting on what data is being sent, who has access to their data, and what they are using their data for (*"If I have to share some sensitive data, where is the data going*


ID	Guideline	CoI	Stakeholder	Problem
G1	AI tools should be open about data practices and share governance with learners.	-	Developer, Researcher, Learner	Access, Trust
G2	AI tools should be accessible and fit into the busy lives of instructors and learners.	Cognitive	Teacher, Learner	Access
G3	AI tools should be informed by learning science and learning theories.	Cognitive	Researcher, Developer	Low Learning
G4	AI tools should be easy to understand and frictionless to use.	Cognitive	Teacher, Learner	Low Engagement, Access
G5	AI tools should be transparent and explainable.	Cognitive	Learner	Trust
G6	AI tools should minimize cognitive load by presenting essential information clearly and simply.	Cognitive	Developer	Low Learning, Low Engagement
G7	AI tools should provide learners with agency over their learning.	Cognitive	Learner	Low Adoption, Low Engagement
G8	AI tools should cultivate learners’ metacognitive and critical thinking skills.	Cognitive	Learner, Researcher	Low Learning
G9	AI tools should support learners’ motivation and cultivate self-efficacy.	Cognitive	Learner	Low Engagement
G10	AI tools should encourage active, constructive, and interactive engagement over passive consumption.	Cognitive	Researcher, Developer	Low Learning, Low Engagement
G11	AI tools should personalize learning based on learners’ knowledge, skills, and abilities.	Cognitive, Teaching	Teacher, Developer	Low Learning
G12	AI tools should offer meaningful challenges that are aligned with learners’ context, goals, and career.	Cognitive, Teaching	Learner	Low Engagement
G13	AI tools should provide accurate, contextualized help when learners need it.	Teaching, Cognitive	Developer	Low Engagement, Trust, Low Learning
G14	AI tools should provide actionable insights to instructors, learners, and researchers.	Teaching, Cognitive	Teacher, Researcher, Learner	Low Engagement, Low Learning
G15	AI tools should integrate with users’ existing educational technology ecosystem.	Teaching	Developer	Low Adoption
G16	AI tools should align with instructional practices and course content.	Teaching	Teacher	Low Adoption, Trust
G17	AI tools should maintain the human touch and supplement rather than replace human instructors.	Social, Teaching	Learner, Teacher, Developer	Low Adoption, Trust
G18	AI tools should scaffold and support learners in developing their social competencies.	Social, Teaching	Learner, Developer	Low Engagement
G19	AI tools should foster social connection and community.	Social	Learner, Developer	Low Adoption, Low Engagement

**Table 2: The 19 design guidelines for AI in adult learning technologies. Each guideline is mapped to the Community of Inquiry dimension it supports, the primary stakeholders it addresses, and the technology problems it seeks to mitigate.**

to? Who has access to it? So that’s completely unknown currently. That information [is] not available anywhere. And moving forward as we are integrating chatGPT with it, it’ll be going to again a third party vendor.”). Higher education students have an implicit trust in their institution to handle data ethically [74], but that trust does not extend to data collected by a third-party educational tool [69]. Clear communication about data use can build trust, reduce user anxiety, and empower students as active stakeholders in how their information is being handled [64].

### 5.1 Cognitive Presence

Sustained engagement in an educational experience allows learners to construct meaning, and is a core component of cognitive presence. Adult students and instructors need technologies that allow them to engage in educational experiences despite the constraints of their daily responsibilities. **Accessibility and flexibility (G2)** emerged as concerns for adults, as both learners and instructors juggle limited time, resources, and devices. Students emphasized the need for tools that accommodate learning differences, are mobile-friendly, and are usable without strong internet access. These features let



**Figure 4: Number of stakeholder comments by design guideline, broken down by group (research team, students, and instructors). Each subplot shows the total volume of comments for each guideline according to the subgroup.**

adults fit learning into their busy lives instead of fitting their life around learning. Instructors similarly reported limited class time and highlighted the value of AI for providing on-demand support, especially outside normal hours (“I might be up at 3:00 AM and if I post a question on the discussion board but you’d usually have to wait until like 8:00 AM 9:00 AM to get a response ‘cause no one’s up

at 3:00 AM”). By aligning with universal design principles to make technology accessible, AI tools can make learning more flexible, inclusive, and sustainable for everyone.

Building technologies with a strong **Theoretical Foundation (G3)** was a concern voiced entirely by the research teams. The development teams behind the AI technologies aimed to ground their

tools in established theories, such as self-determination theory—to maintain motivation through autonomy, competence, and relatedness [21]—or multimedia learning theory—to improve engagement and retention by using multiple representations of content [55]. Although other stakeholder groups mentioned motivation, critical thinking, or active engagement, instructors and students never explicitly discussed the importance of theoretical or empirical evidence for the effectiveness of an educational intervention. Nevertheless, AI technologies should be grounded in evidence to ensure they effectively support desirable outcomes, such as comprehension, motivation, and knowledge construction.

Learning already demands substantial cognitive effort, making it important for learning technologies themselves to be easy to use. Guideline **G4 (Frictionless)** strives to make AI tools intuitive, so students and instructors can focus on learning. Brief, just-in-time onboarding (e.g., tooltips, short videos, examples) should supplement the user interface to clarify expectations and common actions (“*As funny as that might sound, whether there be a video, maybe like an instructional video, just more information on it because I think there was some confusion on how to actually use [the tutor].*”). Adults need brief, contextual, and action-oriented onboarding to support confidence and rapid sense-making [77]. Importantly, this guideline addresses usability-related friction rather than eliminating productive struggle that supports learning. Reducing friction in how a tool is accessed and operated builds confidence and trust in the AI technology. According to the Technology Acceptance Model, perceived ease of use is a primary predictor of users’ confidence, trust, and willingness to adopt new technologies [54].

Students heavily commented on the need for **Transparency (G5)** in AI learning tools. Learners wanted explanations for why they received the guidance they did, especially when AI tools tell them they are wrong. The tools should provide support for learners to construct mental models for how the tool operates (“*I would like to see the inner workings or...what are the data points that essentially are being used to make these recommendations.*”). Prior work on explainable AI in education shows that explanations enable learners to understand how AI systems work, judge the credibility of their outputs, and regulate their own learning [40]. Encouraging self-regulation is essential for adult and self-directed learners [58].

Along with designing tools that are backed by learning science theories, researchers also stressed the importance of designing for **Cognitive Load Management (G6)**. Learners benefit when tools reduce distractions and focus attention on what matters most. Irrelevant or incongruent information places unnecessary demands on working memory and can impair comprehension, problem solving, and memory, particularly for older adults who are more susceptible to distraction [83]. Tools that align design features to instructional content and assessment can reduce extraneous demands on working memory to allow more cognitive resources directed toward learning. At the same time, certain design elements, like interactivity in digital environments, may introduce extraneous demands while still enhancing engagement and learning when they are aligned with the learning goals [73]. By carefully crafting a digital learning environment, AI tools can reduce extraneous demands, improve engagement and motivation, and ultimately increase adoption.

A key characteristic of adults, in contrast to children, is they have greater **Agency (G7)** in their learning [58]. AI tools should offer

choices in how they engage with content and control over how AI supports them. Students consistently voiced a desire for control over how the tool provides support, resources, and content (“*I think just in general, the more options you have for any type of support is always good.*”). Learners should also be provided with guidance to scaffold their ability to make effective choices. This could be done through actionable insights and visualizations (**G14**), contextualized help (**G13**), or developing metacognitive strategies (**G8**). By fostering autonomy and self-direction, AI tools not only reduce barriers to adoption but also strengthen motivation, ownership, and long-term learning outcomes.

Beyond delivering answers, AI should guide learners in using **Metacognitive Strategies (G8)** such as summarization, self-explanation, and reflection to deepen understanding. Tools can also coach students to make thoughtful choices about how they learn—helping them monitor progress, compare approaches, and practice higher-order reasoning. Stakeholders stressed the importance of not just including reflective prompts, but scaffold how to engage in reflection (“*AI is telling you how to reflect, [but is it] giving you the tools to do so?*”). By fostering reflection and critical thinking, AI supports not only immediate task success but also transferable skills for lifelong learning.

**Motivation and Self-efficacy (G9)** were also concepts that stakeholders named as important for cultivating better learning. Researchers wanted to motivate students through in-app messaging, competition, and rewards. Students and teachers wanted to make progress more apparent in the tools but in a way that encourages students to keep trying (“*You know, you watch that little percentage bar not increase or worse, decrease. I could see that being problematic for some people, like causing frustrations or discouraging them.*”). Teachers also voiced the need for step-by-step support so students continue to make progress. Effective AI tools should cultivate learners’ motivation and belief in their ability to succeed.

One of the most commented on topics was the desire **Actively Engage (G10)** students in their learning. Tools designed around the ICAP framework can scaffold inquiry, generative tasks, and peer interaction to strengthen cognitive presence [13]. Learners feel and stay more engaged when prompted to question, explain, and collaborate rather than just consume information (“*I do think that the value is feeling connected and having it feel like you’re actively participating in like class rather than feeling like I’m just watching a bunch of videos and there’s no office hours or there’s no one to talk to.*”). This active engagement helps learners build deeper understanding and retain knowledge more effectively.

## 5.2 Teaching Presence

Within CoI, the organizational structure of the learning experience is described as teaching presence. AI tools can support teachers in their organization by **Integrating (G15)** easily into systems already in use, such as learning management systems. Integration should follow open standards to avoid vendor lock-in, support interoperability, and reduce cost. Effective integration makes the tool feel like a natural extension of the existing ecosystem rather than an added burden.

**Course Alignment (G16)** was a consistent concern from instructors. Teachers repeatedly voiced that they wanted the tools

to support their specific content and practices (*"I don't know about other styles, in my style there will only be 1 choice. Not many choices like in this box."*). Even when teaching the same material, different instructors may structure lessons, sequence activities, or emphasize concepts in distinct ways. Tools with fixed instructional workflows or rigid content structures are unlikely to transfer between different instructors. Without clear alignment to instructional practices and course content, teachers are unlikely to adopt AI tools in their classrooms, regardless of potential benefits or technical quality [59]. For educational AI tools to be successful, they need to earn the trust and support of instructors before the tool is used with students.

### 5.3 Cognitive and Teaching Presence

Many guidelines have the potential to support multiple forms of presence at once. Within the Community of Inquiry framework, the presences are deeply interconnected with one another and this interconnection is necessary to sustain meaningful learning. Some guidelines naturally extend their influence across multiple dimensions of presence rather than being confined to a single category.

**Personalized Learning (G11)** was described as essential for adult learners, who may have incredibly varied prior education and experience. Teachers noted that personalization can help bridge gaps for those balancing work, family, and education, by ensuring that time spent on learning is both efficient and relevant (*"It would really help bridge the gap for a lot of people, particularly [...] people with kids, people with full time jobs... who might not otherwise be able to get instruction in person."*). By engaging the students through tailored learning experiences, this guideline seeks to increase cognitive and teaching presence. Personalized learning pathways are essential for keeping adult learners engaged, motivated, and making continual progress [70].

Learners also want to make sure that learning is **Aligned with Future Goals (G12)**. Knowles [44] says that adults are problem-oriented, and need to know how what they learn supports why they learn. Students want to see how the topics they are learning relate to future careers. AI tools should make clear how the content supports students' learning goals so adult learners stay engaged and motivated by giving students meaningful challenges that support learners' careers. By situating learning within contexts that matter to adults, AI tools can support persistence and engagement. This guideline can enhance cognitive presence by connecting the content to authentic applications, and teaching presence by aligning the educational experience with the learner's long-term goals.

Learners stressed the importance of receiving relevant help when they need it. **Contextualized Help (G13)** states that AI tools should recognize when learners need help and give them accurate, proactive instructions. Students voiced frustration with the help methods currently available in AI tools, while also raising concerns about trusting the accuracy of an AI's response (*"the big [problem] is you're given a wrong answer, but you don't realize it"*). AI tools must give accurate, contextualized answers without errors or hallucinations or user trust quickly erodes. This guideline supports increasing cognitive presence by engaging with the student in the learning process through feedback loops, and teaching presence by providing clear instructional guidance.

**Actionable Insights (G14)** within AI technologies can increase teaching presence by helping instructors design instruction and increase cognitive presence by helping learners self-monitor their progress. Teachers wanted to see who was using the tools, what activities they engaged in, and how the students progressed over time. Teachers then wanted to use this information to improve their instruction and content sequence. Students also voiced that data analytics help them keep pace in the class and reflect on their progress (*"For me that helps me like just getting a visual representation of where I am and where I need to be and what I'm missing. That helps me a lot actually."*). AI tools should move beyond presenting raw data to providing actionable insights that inform instructional decisions and support learner reflection [86].

### 5.4 Social Presence

Within the Community of Inquiry framework, social presence allows learners to project their authentic selves into the learning environment to interact with others in a meaningful way. Students cited social connection with other humans as something they do not want to lose through interacting with AI tools. **Fostering community (G19)** should be a goal of AI tools, which could be realized by helping learners find peers with shared goals, interests, or contexts. Adult learners want to make friends, not just for social connection, but also because learning communities are an important part of education. Learning on your own can be disheartening, (*"If I'm struggling with the class, it feels like I'm struggling alone."*) but AI tools can help to offset those feelings both by providing content support, and connecting learners together.

### 5.5 Teaching and Social Presence

One of the most talked about concerns for adult learners was social connection. The following guidelines seek to increase social presence through lowering the barrier to social interaction and preserving interpersonal connections. They also can strengthen teaching presence by centering the instructor's responsibility for course design and intentionally creating ways for students to learn social skills.

Adult learners consistently emphasized the importance of human presence in their educational experiences, expressing greater trust in guidance from humans than from AI. **Preserving the Human Touch (G17)** highlights how learners value feedback, explanations, and materials grounded in human judgment rather than automated responses. Students worried that AI-driven instruction removes the interpersonal elements of learning, particularly when individual needs are involved. Others described a broader concern that increasing reliance on technology makes education feel transactional and diminishes opportunities for meaningful connection with instructors. While AI tools may increase access to information and efficiency, students cautioned that overdependence could erode patience, empathy, and tolerance for frustration (*"I have the slippery slope idea that if people have these robot servants at their fingertips at all times, that they will have zero tolerance for frustration."*). To build trust and engagement, AI should supplement rather than replace human involvement to preserve the relational dimensions of learning that students find most meaningful.

Making genuine connections can be challenging, and soft skills like social interaction are skills that AI tools can help students learn. **Scaffolding Social Competencies (G18)** can help learners make friends while teaching them how to reach out to people without the assistance of an AI tool. Initiating conversations with strangers can be uncomfortable, and asking for help in a large, public forum can feel intimidating. Consulting an AI tool with questions could lower the barrier to receiving help since students may feel that classmates or instructors will judge them. At the same time, some students worried how over-reliance on AI may impact their real-world social skills (*“If you over-rely on the tool, like, if you just purely use it as a source of meeting new people, it can kind of stand in the way of you meeting people in real life.”*). AI tools should not only connect students to build a learning community, but also support their development of social competencies so they will not need the tool to make future connections.

## 6 Application of Design Guidelines

### 6.1 Heuristic Evaluation

Heuristic evaluation lets researchers systematically analyze interactive systems in order to surface strengths, limitations, and design trade-offs in a structured and comparative manner [67]. Here, we operationalize our design guidelines by conducting a heuristic evaluation of the AI-ALOE technologies described in Section 3.2, using them as an analytic lens to examine how well these systems support adult learning needs and identifying opportunities for improvement. Because these technologies were developed without an explicit set of adult-learning-specific design guidelines, they provide a natural test case for illustrating how the guidelines can be used to support evaluation, ideation, and reflection.

**6.1.1 Technology Coding.** To carry out this evaluation, we recruited four independent coders to apply the full set of design guidelines to each of the AI-ALOE technologies. Coders worked independently and evaluated multiple systems, with each technology reviewed by two coders. Coders relied on a combination of demo videos and peer-reviewed publications describing the technologies to understand user interactions and underlying design intentions. Where aspects of a system were ambiguous or not immediately evident from demonstrations alone, coders consulted the associated literature to clarify system functionality.

Coders used a structured rubric to guide their evaluations. The rubric consisted of each guideline alongside its description, framing the guideline as an evaluative statement applied to a specific technology (e.g., “[This tool] is easy to understand and frictionless to use”). For each guideline, coders rated the degree to which there was evidence that the technology applied the guideline on a five-point Likert scale, ranging from strong evidence that the guideline was not met to strong evidence that it was satisfied. Coders were also required to provide a written justification explaining their assessment. An excerpt from the heuristic evaluation rubric is provided in Table 3.

To resolve differences in ratings, coders participated in structured meetings to discuss their evaluation and how they applied each guideline. During these sessions, coders explained the reasoning behind their scores, clarified how they interpreted the intent and

scope of the guidelines, and examined points of disagreement. This collaborative process allowed coders to refine their judgments and reach greater consistency in how guidelines were applied across technologies.

After consensus was reached, we computed the average rating for each technology-guideline pair and shared these results with the research and development teams responsible for each system. Teams were invited to review the evaluations and provide additional documentation if they believed aspects of their technology had been mischaracterized. Coders reviewed the additional materials, revised ratings where appropriate, and reconvened to discuss any changes. This process resulted in a strong overall agreement across coders (intra-class correlation coefficient = 0.80) when determining whether a guideline was satisfied. The aggregated results of this evaluation are summarized in Figure 5.

**6.1.2 Results.** Across the evaluated technologies, several guidelines were consistently rated as satisfied. The highest average scores were associated with usability (**G4**) and integration with existing educational technology ecosystems (**G15**). Most tools were compliant with the Learning Tools Interoperability (LTI) standard and did not require substantial additional setup or configuration when used. The technologies also use familiar web-based interfaces, like chat interfaces, textboxes, and drag-and-drop components, to allow users to onboard themselves to the technology quickly.

A second cluster of highly rated guidelines reflected cognitive and instructional design theories: **G3** (informed by learning science), **G6** (minimizing cognitive load), and **G10** (active, constructive, and interactive engagement). Interfaces were designed to reduce distractions and focus attention on what matters most. Learning activities frequently required active participation, exploration, and reflection, rather than passive content consumption. As a result, the evaluated tools were consistently rated as supporting interactions grounded in well-established instructional principles.

The technologies also scored highly on **G5** (transparent and explainable). In particular, systems often provided explicit explanations for feedback and guidance, especially in response to incorrect learner actions. These explanations helped clarify why particular forms of feedback were presented. Across technologies, this form of explainability was more prevalent in learner-facing feedback mechanisms than in system-level descriptions of underlying models or processes.

In contrast, several guidelines were consistently rated as less well supported. The lowest average scores were associated with social interaction and community-related guidelines, including **G18** (scaffolding social competencies) and **G19** (fostering community). For most technologies, opportunities for peer-to-peer interaction were limited, and explicit support for developing social skills was either minimal or absent. These patterns were consistent across tools, with few systems incorporating features explicitly designed to support sustained social engagement.

Another guideline that received low satisfaction scores across technologies was **G1**, which concerns transparency around data practices. Coders noted that information about data collection, storage, access, and use was often not readily visible to users through system interfaces. As a result, it was difficult to determine how user

ID	Guideline	Sample Score	Rubric Description
G1	[This tool] is open about data practices and share governance with learners.	3	Evaluate whether data collection, storage, use, and access are clearly communicated to users. Higher scores reflect meaningful user control (e.g., ability to review, manage, or delete data) and understandable disclosures. Lower scores reflect opacity, unclear retention or access, or limited user agency over personal information.
G2	[This tool] is accessible and fit into the busy lives of instructors and learners.	2	Evaluate whether the tool can be used quickly and reliably across contexts (work/home/school), times (including nights/weekends), and devices, with reasonable cost and effort. Higher scores reflect low barriers to access and flexible modes of engagement. Lower scores reflect restricted access contexts, or poor support for diverse needs.
G3	[This tool] is informed by learning science and learning theories.	4	Evaluate whether the design rationale and core interactions are grounded in established learning theories. Higher scores reflect explicit theoretical grounding and design choices that apply theory to support comprehension, motivation, and knowledge construction. Lower scores reflect absent or inconsistent theoretical justification, or features that may conflict with relevant theory.

**Table 3: Excerpt from the heuristic evaluation rubric, illustrating how AI tools could be evaluated against the guidelines for reflection and ideation.**

data were being handled or what degree of control users had over their information.

**G11** (personalization) was another area with relatively low satisfaction. While many technologies incorporated surface-level personalization – such as adapting content examples, responding to learner inputs, or providing regular knowledge checks – fewer systems used learner data for adaptations such as task sequencing or calibrating difficulty. Personalization was therefore present in limited forms but was not consistently reflected in deeper structural adaptation across learning activities.

Satisfaction scores varied across individual technologies, revealing distinct strengths and limitations. For example, SAMI, as a social agent designed to connect learners, was the only system that consistently satisfied the social presence guidelines (**G18**, **G19**). At the same time, SAMI received lower scores on guidelines related to metacognitive support (**G8**), personalization (**G11**), and instructional alignment (**G16**) since these fall outside its primary function. Jill Watson, as a conversational virtual TA, engages and empowers students to ask questions and quiz themselves, building motivation and self-efficacy, so it showed strong support for learner motivation and engagement (**G9**), while receiving lower ratings for grounding in learning science theories (**G3**). VERA’s ability to build scientific concepts through experimentation and scientific investigation earned a high rating on learning science-related guidelines (**G3**, **G7**, **G8**), but the tool received lower scores on motivation-oriented guidelines (**G9**, **G11**).

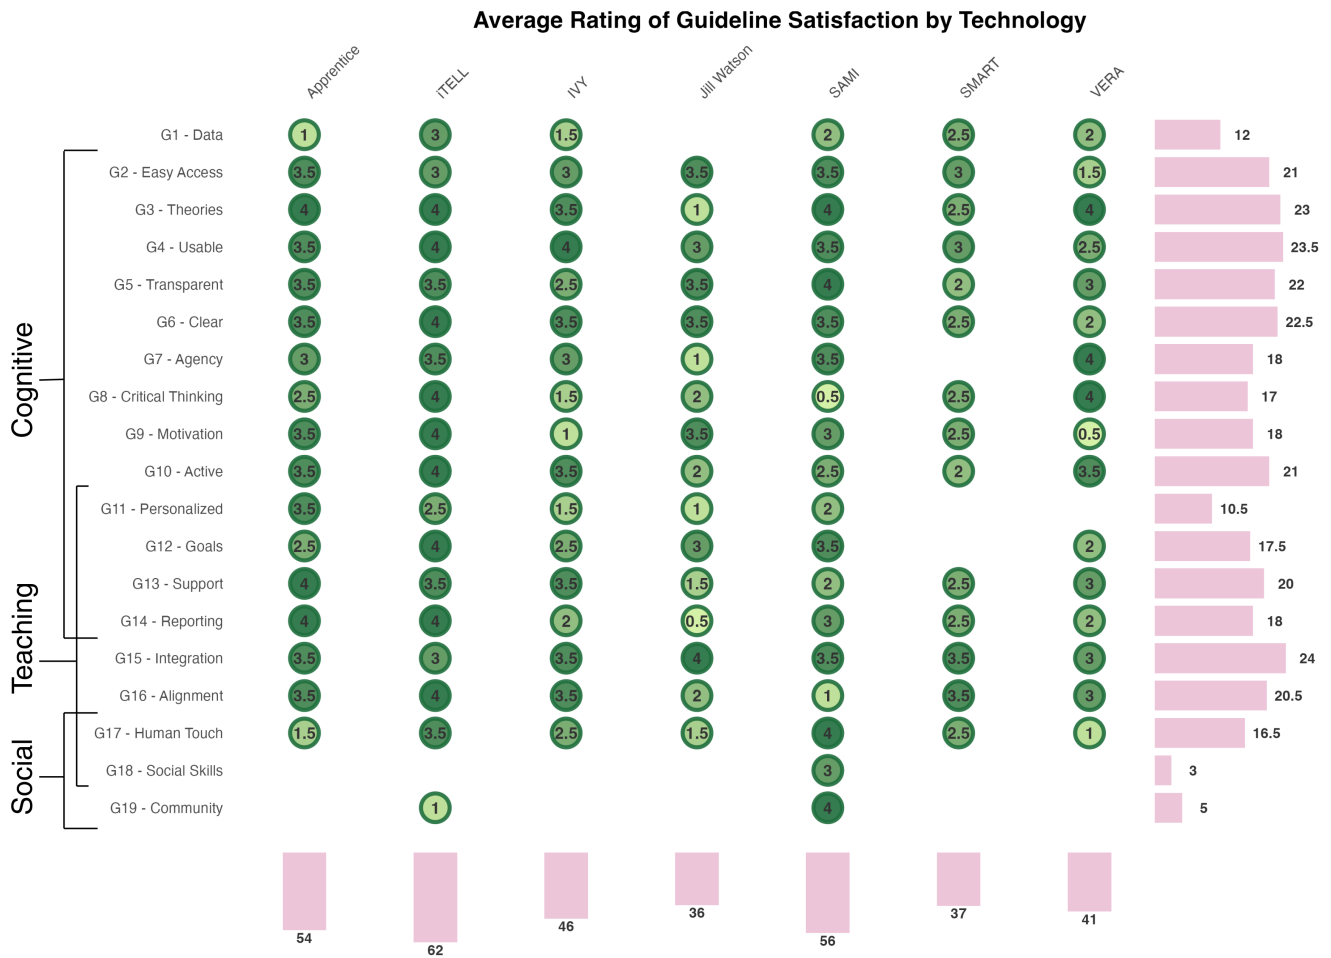
Overall, no single technology satisfied all 19 guidelines. Instead, the heuristic evaluation revealed variation in guideline coverage that aligned with the functional focus of each system. Technologies designed to prioritize particular instructional, cognitive, or

engagement-related goals tended to support corresponding guidelines more strongly, and providing minimal and limited support for others. When considered collectively, however, the broader ecosystem of AI-ALOE technologies provided coverage across the full set of guidelines. From this perspective, lower satisfaction scores on individual guidelines reflect trade-offs in scoped design decisions rather than an absence of support at the ecosystem level.

## 6.2 Guideline Explorer

We developed the Guideline Explorer (Figure 6) as an interactive tool that provides direct access to the ideation statements—the first-person expressions derived from the blue sticky notes generated during the affinity diagramming process described in Section 4.2. These statements capture recurring needs articulated by learners, instructors, administrators, and developers across deployments. To make the data navigable, we annotated each ideation statement with metadata including its associated guideline, Community of Inquiry dimension, affected stakeholder group, and problem type. The Guideline Explorer enables designers to browse, search, and filter ideation statements along these dimensions. We envision designers first using the heuristic evaluation rubric to identify opportunities for growth within a technology, and then turning to the explorer to examine associated stakeholder needs.

A contribution of the explorer is that it makes qualitative design research directly accessible and actionable beyond our specific research context. When designing or refining AI-powered education technologies, developers can read stakeholder statements from actual adult learning contexts to help determine if and how they might incorporate specific guidelines in their technology. By engaging with ideation statements rather than specific, prescriptive solutions,



**Figure 5: Average satisfaction ratings for each design guideline across the research institute’s technologies. Green bubbles show average ratings by two coders on each technology from 0 (there is strong evidence the technology does not satisfy the guideline) to 4 (there is strong evidence the technology does satisfy the guideline). If the guideline was rated as 0 by both coders, the bubble is missing to highlight a potential improvement area for the technology. Pink bars on the bottom show the total rating across all technologies, and the pink bars on the right show the total rating across each guideline. Guidelines vary in satisfaction, with the strongest support for integration, usability, learning theories, cognitive load management, and transparency (G15, G4, G3, G6, G5), and the lowest support for social skills, community, personalization, and open data practices (G18, G19, G11, G1).**

designers can brainstorm and evaluate potential responses that are sensitive to their own context. We should note that although some of the ideation statements explicitly mention AI, one should not assume that AI-based solutions are always appropriate. Instead, it surfaces stakeholder concerns that may be addressed through a range of means.

Next, we present two cases detailing how the explorer might be used.

**6.2.1 Case 1.** Imagine a team used the heuristic evaluation and found that its technology was rated low on guideline **G14** – *AI tools should provide actionable insights to instructors, learners, and researchers*. The research team knows that they are collecting lots of data, but the main use of that data is learning engineering and

their own analyses. The AI tool provides actionable insights to researchers, but not learners or instructors. They open up the Explorer and investigate the **G14** ideation statements:

- **Instructors and Researchers:** I want to know who are using the tools, who did what, and to track their progress.
- **Instructors and Students:** I want detailed learning analytics (through visualizations) so I can make better instructional decisions.
- **Researchers:** I want the tools to collect data and use it to facilitate teacher and tool personalization.
- **Researchers:** Learning tools should be designed for continuous improvements and learning engineering through historical data.

# AI Technology Design Guidelines Explorer

An interactive tool to search and explore design guidelines for AI technologies that support Adult Learners.

**Search Guidelines**

🔍 Keywords or "exact phrase" All Text Ideation Guideline

**Filters** Clear All Filters

Group By: COI Dimension Stakeholders Problems COI Dimensions

## Cognitive

**Ideation Statement**

*"The AI tool has a big learning curve and I need an introduction / tutorial on what it can do and how to use it"*

Stakeholder: Teacher Learner

Problem: Low Engagement Access

CoI: Cognitive

**Ideation Statement**

*"I need an AI tool that works without errors and is easy for me to understand quickly, or I won't use it (because my time is better spent elsewhere)"*

Stakeholder: Teacher Learner

Problem: Low Engagement Access

CoI: Cognitive

**Ideation Statement**

*"I want to design AI tools that are intuitive, easy to use, and promote natural interactions to avoid confusion over the interface and expectations of the tool"*

Stakeholder: Teacher Learner

Problem: Low Engagement Access

CoI: Cognitive

**Figure 6: The Guideline Explorer interface for browsing and filtering ideation statements associated with each design guideline. The explorer allows designers to search statements and filter them by guideline, stakeholder, Community of Inquiry (CoI) dimension, and problem type to support stakeholder-centered ideation.**

The team can now brainstorm ways to meet the stakeholder needs around actionable insights. Perhaps the data is not yet in a consumable format for stakeholders outside the research team, so they focus on transforming the data into something more interpretable like a table that provides an overview of class actions to meet the first ideation statement. Or, they could brainstorm ways to personalize the tool based on data so that each user has an experience tailored to their needs, meeting the third ideation statement. Maybe the team decides to build an instructor dashboard to meet the first statement, detailing who, what, and when, or a more detailed visualization of class-level misconceptions to help instructors structure in-class time around student needs. Because self-directed adult students also play a role in instructional decision making, perhaps the team proposes a student-facing dashboard to make student progress visible and help students monitor and reflect on their learning trajectory.

Or, like SMART, the team could implement all of these ideas (Figure 7). After reviewing the user interview transcripts, the SMART

team came to similar conclusions about how to improve their tool. If resources like the guidelines and the explorer had been available earlier, the team could have used them to more quickly identify promising areas for improvement. This does not suggest that teams should not engage with users; rather, it demonstrates how these resources can efficiently synthesize prior research into actionable design directions.

**6.2.2 Case 2.** Alternatively, assume a technology team is struggling with low adoption. They developed their technology with the input of users and stakeholders from one university, but struggle trying to move their technology to another. Their technology is embedded in several classes and they see many initial launches, but very few users continue to interact with the tool beyond the first few clicks. The team decides to investigate stakeholder concerns related to low adoption using the explorer. They find the following ideation statements:



**Figure 7: Example interfaces from SMART illustrating how learner data can be transformed into actionable insights for different stakeholders, supporting Guideline G14. Three panels show new changes made in response to learner and instructor needs. Panel A shows a timeline-based visualization of learner activity and revisions. Colored blocks represent different revisions or learning events over time, and hovering over a block reveals detailed metrics such as time spent, number of words, and number of concepts for fine-grained analysis of learning behaviors. Panel B presents an instructor-facing table summarizing student performance and progress. This view supports instructors in monitoring class progress, identifying students who may need support, and making instructional decisions based on summarized learning analytics. Panel C displays benchmark and individual statistics for reflection and comparison. Metrics such as total duration, number of revisions, and achieved level are shown for both the class average and the individual student, supporting reflection and self-monitoring.**

- **Instructors:** I won't use the AI if it isn't aligned with my teaching approach and style.
- **Instructors and Students:** I want ways to appropriately integrate AI with my other tools and to effectively manage AI tool notifications.
- **Instructors and Students:** I want options/choices for how I engage with content, as well as access to supplementary content.
- **Students and Researchers:** I want AI to provide explanations and sources from human experts so I can believe what the AI says is correct.

The team can now brainstorm ways to improve their tool toward greater adoption. The first point is a statement from instructors, so the team could investigate the instructional practices of the new school. How well does their tool align with and reinforce the existing instructional practices (G16)? The second point is from both instructors and students about integrating into existing technology ecosystems (G15). Is the AI tool connected to the other tools currently in use? Are the notifications from all tools easily consolidated? Can students and instructors efficiently navigate between

each technology? The third point is a concern raised by students asking for agency over their learning (G7). Does the tool provide options for students or is there only one way to engage? Are students locked into a prescribed path or can they make meaningful choices about their learning trajectory? Does the tool provide optional, supplementary content and resources? The fourth point is also a concern from students regarding transparency and explanations within AI tools (G5). Does the AI tool link back to human-crafted source material to justify its explanations? Does the AI tool build trust by proving that its content is correct?

These questions show how the explorer can also support designers in looking across guidelines to uncover potential root causes of design challenges such as low adoption. Instead of treating a lack of adoption as a single problem, the team can trace it to interacting issues of instructional alignment (G16), ecosystem integration (G15), learner agency (G7), and transparency and trust (G5). Now the design team can make concrete changes to their tool, such as modifying instructional workflows, enabling integration with existing platforms, offering multiple learning paths, or adding citations and explanations to AI responses. This kind of synthesis exposes relationships among guidelines that would be difficult to see when

they are considered independently. The explorer can help designers move from vague problem statements to coherent hypotheses about why a tool is not working through cross-guideline analysis.

## 7 Discussion

The guidelines presented in this paper are grounded in the perspectives, concerns, and desires of key stakeholders: learners, instructors, and researchers. They reflect priorities from participants' lived experiences with AI-enhanced learning technologies. Even when individual guidelines align with principles that are already discussed in prior literature, empirically demonstrating their significance for this specific population of adult learners constitutes a meaningful contribution. Many of the principles, like transparency, accessibility, and cognitive load management, are rooted in foundational learning science principles and are applicable across age groups and educational settings. Their inclusion here demonstrates their relevance in practice with this under-studied population. Validating these principles in the context of adult learning confirms their importance while grounding them in the realities of this population.

At the same time, the guidelines surfaced many concerns that are specific to the adult learner. Adult learners have many responsibilities competing for their time, so they must divide their attention. AI educational tools should be easy to access wherever, whenever (G2), be easy to use and understand (G4), and designed to minimize extraneous cognitive load (G6) so that learning can fit into adults' limited and fragmented time. Adults are also more likely to be concerned with data governance (G1) than child learners. In addition to minimal time, adults are career-oriented, so learning experiences should be purposeful and connected to the learner's real-world goals (G12). Adults have varied backgrounds, and bring a wealth of prior knowledge and experiences into the classroom. AI tools should provide personalized pathways based on the adult learner's background (G11), provide timely and contextualized support when needed (G13), and give learners meaningful agency over how they engage with and direct their learning (G7). These guidelines operationalize broader design principles in ways that are directly responsive to the needs of adult learners.

We view these adult-focused guidelines through the lens of universal design: design choices that address the needs of adult learners often improve the learning experience for all users. Universal Design [25] is a framework originally developed for the built environment, but these principles have since been applied broadly to communication, technology, and education, based on the insight that design choices that address the needs of users with diverse abilities, constraints, or preferences often improve the experience for everyone. Universal Design emphasizes flexibility in use, guiding designers to "Provide choice in methods of use" like our G7, giving learners agency over their learning, and "Provide adaptability to the user's pace" like our G11, where AI tools should provide individualized pathways for learning based on learners' knowledge, skills, and abilities. The framework also emphasizes presenting information in easy-to-understand, accessible ways ("Provide compatibility with a variety of devices." and "Eliminate unnecessary complexity.") similar to G2 and G6, which strive to make AI tools accessible whenever, wherever and present information clearly and simply.

Building on this foundation, Universal Design for Learning (UDL) applies these principles specifically to educational contexts through three core principles: multiple means of representation (how information is presented), action and expression (how learners demonstrate knowledge), and engagement (how learners are motivated and supported) [11]. Features such as flexible access and personalized support address representation by ensuring information reaches learners in forms they can perceive and process. Learner choice, agency, and goal-oriented experiences support action and expression by giving learners control over how they demonstrate and apply knowledge. And motivational scaffolds, social connection, and metacognitive development exemplify UDL's engagement principle, which emphasizes sustaining effort, welcoming identities, and cultivating emotional capacity. Examining the unique needs of adults surfaces requirements around flexibility, relevance, and efficiency that would benefit all learners with limited time, varied backgrounds, or different learning preferences across educational contexts. In this sense, focusing on the constraints and priorities of adult learners does not narrow the applicability of these guidelines. Rather, it strengthens their general value by promoting more inclusive, usable, and effective AI-enabled learning technologies for a broad range of learners.

One issue that came up during our discussions with AI-ALOE researchers was whether there should be an additional guideline concerning the ethical use of AI. Although the guidelines do not explicitly frame themselves as ethical principles, we believe several provide concrete and actionable ways to build AI technologies with ethics in mind. For example, transparency and explanations (G5) is an actionable response to concerns about hallucinations, misinformation, and other limitations of AI-generated content. Data governance (G1) translates into making user privacy and consent visible and giving users more ownership of their data. Guidelines around learner agency (G7), contextualized support (G13), instructional alignment (G16), and preserving the social elements of learning (G17-G19) reflect ethical considerations around autonomy, accountability, and the risk of losing human connections. The guidelines offer an actionable path for embedding ethical considerations into AI-supported learning environments.

An important pattern that emerged from our analysis is the unequal distribution of design guidelines and empirical support across the Community of Inquiry dimensions. While there is a strong emphasis on cognitive and teaching presence, there is comparatively little attention to social presence. Of the 19 guidelines, only three (G17-G19) explicitly address the social dimensions of learning, and these guidelines received the lowest satisfaction ratings across the AI-ALOE technologies. Even when AI-ALOE tools included features to support social presence, they were often peripheral rather than a central function, with SAMI being a notable exception. This imbalance suggests that AI-supported adult learning tools are relatively mature in supporting individual cognition, usability, and instructional alignment, but underdeveloped in fostering social connection, building community, and developing interpersonal skills.

In addition, concerns about social presence and the erosion of human elements in online learning environments were primarily expressed by students (see Figure 4). Learners repeatedly emphasized the importance of interpersonal interaction and human judgment in learning, apprehensive that AI systems could reduce opportunities

for meaningful connection with instructors and peers. Participants also worried that AI tutors and automated feedback could hide contextual factors such as learning disabilities, emotional states, or personal struggles that humans are better equipped to recognize. Learners also expressed concerns about blindly accepting information from an AI; e.g., in case the virtual TA reports untrue information about assignment deadlines or an AI tutor provides incorrect remediation.

Adult learners have clearly articulated that learning is not solely a cognitive activity. Relationships, community, and connection are essential components of a learning environment. Students continue to prefer learning from human experts, as AI systems have not yet demonstrated the trustworthiness, accountability, contextual understanding, and human presence required to be accepted as instructional authorities. Future AI learning technologies should move beyond performance-oriented goals and more intentionally design for social presence and shared inquiry.

In addition to gaps in social presence, our analysis reveals potential tensions within the guidelines that reflect differing priorities across stakeholder groups. One such tension exists between the guidelines that AI tools should be easy to understand and frictionless to use (G4) while providing quick and proactive help (G13), and the guidelines that emphasize meaningful challenge (G8), metacognitive development (G10), and active engagement (G12). Learners and instructors wanted to reduce usability friction, like confusing interfaces, unclear expectations, or excessive onboarding demands, as it makes the tool feel efficient and supports ongoing work. Delays or unclear responses were frustrating to users, especially if students wanted to get the answer and move on. However, the desire for immediacy may conflict with opportunities for learners to engage in deeper thinking. Quickly providing answers can reduce the time and cognitive effort students spend generating and evaluating their own responses, potentially limiting reflection and understanding. Rather than optimizing for speed or responsiveness, AI tools may need to balance timely support with mechanisms that encourage learners to pause, consider their own reasoning, and remain cognitively engaged. When intentionally designed, friction can promote slower, more mindful engagement and improve users' understanding of task goals and overall satisfaction [57]. Well-designed AI tools can preserve productive struggle through meaningful challenges and reflective prompts while minimizing operational friction that distracts from learning goals.

At the same time, researchers were the only stakeholder group to emphasize that AI learning technologies should be grounded in learning theory and empirical evidence of effectiveness. The competing desires between students, instructors, and researchers suggests a risk that systems optimized for immediacy and ease of use may be preferred by learners and instructors, even when those designs may undermine learning processes that rely on productive struggle and sustained cognitive effort. When ease of use is a significant deciding factor for adoption, prioritizing convenience over reflection and sense-making becomes a danger for learning technologies. Future AI learning technologies should make the value of productive struggle more visible to users, so they better understand the benefits of slowing down and engaging in deliberate practice.

Another tension arises between learner agency (G7), personalization (G11) and instructional alignment (G16). Adult learning theory

emphasizes autonomy, self-direction, and choice, and learners frequently expressed a desire for control over pacing, support, and learning pathways. AI systems are well-suited to providing many options for students through personalization, optional scaffolds, and adaptive pathways. However, instructors voiced parallel concerns about alignment with course goals and assessments. Learner agency in adult education does not imply total freedom, but rather meaningful choice within a structured instructional context. When AI tools make instructional goals, expectations, and rationales visible, they enable learners to exercise agency in relation to those goals. Instructional alignment makes agency productive by orienting learner's choices around pacing, strategy, and support toward course objectives. AI tools can translate instructor intent to students in a way that preserves meaningful learner autonomy.

The guidelines have already informed concrete revisions across multiple systems, including Apprentice, iTELL, Jill Watson and SMART. Jill Watson, iTELL, and Apprentice were revised to make all features compatible on mobile devices, making the tool accessible from many locations and times (G2). The iTELL development team enabled universal integration with industry learning management systems using SCORM broker files (G15), and added AI-generated feedback for constructed responses to ensure that support is contextualized and relevant (G13). Apprentice also expanded its feedback mechanisms to offer deeper hints, supplementary videos and texts, and problem-specific tutorials (G13). In addition, Apprentice developers collaborated with instructors to create multiple versions of the same tutored content to ensure alignment with instructional practices (G16). SMART also incorporated new dashboard features to learners and instructors, as described in Section 6.2.1. These examples demonstrate how stakeholder-informed guidelines can be translated into concrete features across platforms.

Our heuristic evaluation revealed significant variation in how individual technologies align with the design guidelines. Systems tended to strongly support guidelines that were closely aligned with their primary instructional purpose and functional scope. This variation revealed clear trade-offs in design focus. Technologies that emphasized depth in a particular dimension of adult learning often did so by deprioritizing features associated with other dimensions. As a result, no individual system demonstrated strong support across all 19 guidelines. When considered collectively, however, the AI-ALOE technologies provided strong coverage across the full set of guidelines.

These results suggest that the guidelines can support coordination across an ecosystem of complementary technologies, rather than requirements for any single system. Using the Community of Inquiry framing, AI technologies may choose to prioritize cognitive, teaching, or social presence, and focus on the set of guidelines that align with the goals of the tool, while leaving others unsatisfied. However, there may be some guidelines that are 'non-negotiable' to users and may impact the tools' usefulness and broader adoption. Guidelines around usability (G4), accessibility (G2), and instructional alignment (G15) may need to be satisfied by every tool within an educational ecosystem. A tool that is confusing, inaccessible, and does not support instruction undermines the learning experience, and no other tool can compensate for that failure. Other guidelines like G11 and G19 could be distributed across a set of complementary tools, allowing for a deeper personalization as learners

navigate the ecosystem. Building a singular tool that satisfies all 19 guidelines would be a difficult task, but a thoughtfully coordinated ecosystem can achieve what no single tool can alone. In this framing, the guidelines can help designers, instructors, and institutions reason about priorities, trade-offs, and how multiple tools can work together to support the diverse cognitive, teaching, and social needs of adult learners.

## 8 Limitations and Future Work

Our guidelines should be interpreted with respect to the context from which we collected data, that is, adult learning environments. While these tools represent a wide variety of functions, they may not capture the full spectrum of educational technologies in use across different institutions and contexts. Though we have demonstrated the practicality and applicability of the guidelines to produce meaningful change in AI technologies for adult learners, it is unclear if and how the guidelines can be applied beyond adult and online learning contexts. Some of the guidelines may generalize across different age groups and modalities, while others, such as scaffolding social competencies and supporting career-oriented goals, may be specific to adult learners balancing professional, personal, and educational responsibilities.

In addition, our analysis relied on thematic coding and heuristic evaluation. Although we used reflexive thematic analysis and consensus-building sessions to strengthen reliability, the judgments necessarily reflect researcher and coder interpretation. While the guidelines capture stakeholder concerns and priorities, further empirical validation is needed to assess their impact on learning, engagement, and adoption in practice.

Future work should extend these guidelines through empirical validation, examining how these guidelines relate to learning outcomes, motivation, and technology use in authentic instructional settings. Cross-institutional studies, such as workforce training, community colleges, or international contexts, could test whether the guidelines hold in settings with different learner demographics, institutional structures, and resource constraints. Such work would help determine which guidelines are broadly applicable and which are more context-specific.

## 9 Conclusion

Through a reflexive thematic analysis of stakeholder feedback, we derived 19 design guidelines that articulate recurring needs, challenges, and values across adult learners, instructors, and researchers. Organized through the Community of Inquiry framework, these guidelines emphasize the cognitive, teaching, and social dimensions of adult learning while also addressing foundational concerns such as transparency, data governance, and accessibility. They provide a structured vocabulary for reasoning about AI-supported learning beyond technical performance, highlighting issues of agency, trust, alignment with instructional practice, and the preservation of human connection.

We demonstrated the practical utility of these guidelines through a heuristic evaluation of the AI-ALOE technologies. This evaluation revealed both strengths, such as usability, theoretical grounding, and explainability, and areas for growth, including support for social connection, personalization, and open data practices. Rather than

functioning as a checklist, the guidelines are intended to support reflective design: making trade-offs visible, clarifying whose needs are prioritized, and enabling teams to reason about how individual tools contribute to a broader learning ecosystem.

To support this process, we introduced the Guideline Explorer as a companion that links each guideline to underlying stakeholder concerns and ideation statements. The Explorer enables designers to move from abstract principles to concrete design opportunities by making the guidelines and supporting affinity diagram easily navigable and searchable. The guidelines, heuristic rubric, and Explorer offer an approach for evaluating existing systems, informing new designs, and facilitating discussion among interdisciplinary teams.

The guidelines are derived from technologies developed within the AI-ALOE research institute and reflect the contexts, populations, and design priorities of those deployments. While the setting spans diverse disciplines and institutions, further research is needed to examine how the guidelines generalize to other adult learning environments, including workplace training, informal learning communities, and industry-developed platforms. As generative AI and adaptive systems become increasingly prevalent in education, there is a growing need for empirically grounded frameworks that help align technological innovation with the realities of adult learning. AI is not the solution to all educational technology problems, but analyzing the concerns voiced by users can ground ideation and problem-solving in stakeholder needs.

Ultimately, our work aims to support lifelong learning by supporting the design of AI technologies that are grounded in adult learners' needs, resources, challenges, and lived contexts. Designing for adult education requires more than adapting tools built for younger learners. It demands sustained attention to the social, cognitive, and institutional conditions in which adults learn. By designing learning tools based on adult learners' needs, resources, challenges, and lifestyles, researchers and technology designers can improve engagement, social connections, and learning outcomes for adult learners specifically. The guidelines presented in this paper lay the foundation for designing AI that respects and supports the complex lives of adult learners.

## Acknowledgments

The material in this paper is based upon work supported by the National Science Foundation under Grant No. 2247790 and Grant No. 2112532. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

## References

- [1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.
- [2] Sungeun An, Robert Bates, Jen Hammock, Spencer Rugaber, Emily Weigel, and Ashok Goel. 2020. Scientific modeling using large scale knowledge. In *International Conference on Artificial Intelligence in Education*. Springer, 20–24.
- [3] Sungeun An, Jennifer Hammock, and Ashok Goel. 2025. How Online Learners Engage in Self-Directed Modeling: A Behavioral Analysis. *International Journal of Artificial Intelligence in Education* (2025), 1–28.

- [4] Yoojin Bae, Jinho Kim, Andrea Davis, and Min Kyu Kim. 2024. A Study on AI-Augmented Concept Learning: Impact on Learner Perceptions and Outcomes in STEM Education. In *Proceedings of the 18th International Conference of the Learning Sciences-ICLS 2024*, pp. 1450–1453. International Society of the Learning Sciences.
- [5] Matthew L Bernacki, Meghan J Greene, and Nikki G Lobjcowski. 2021. A systematic review of research on personalized learning: Personalized by whom, to what, how, and for what purpose (s)? *Educational Psychology Review* 33, 4 (2021), 1675–1715.
- [6] Melissa Bond, Olaf Zawacki-Richter, and Mark Nichols. 2019. Revisiting five decades of educational technology research: A content and authorship analysis of the British Journal of Educational Technology. *British journal of educational technology* 50, 1 (2019), 12–63.
- [7] John D Bransford, Ann L Brown, Rodney R Cocking, et al. 2000. *How people learn*. Vol. 11. Washington, DC: National academy press.
- [8] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [9] Virginia Braun and Victoria Clarke. 2019. Reflecting on reflexive thematic analysis. *Qualitative research in sport, exercise and health* 11, 4 (2019), 589–597.
- [10] Stephen Buckley, John Kos, Rahul Dass, Cathy Teng, Kenneth Eaton, Sareen Zhang, and Ashok Goel. 2025. A Personalized AI Coach to Assist in Self-Directed Learning. In *Proceedings of the Innovation and Responsibility in AI-Supported Education Workshop (Proceedings of Machine Learning Research, Vol. 273)*, Zichao Wang, Simon Woodhead, Muktha Ananda, Debshila Basu Mallick, James Sharpnack, and Jill Burstein (Eds.). PMLR, 221–229. <https://proceedings.mlr.press/v273/buckley25a.html>
- [11] CAST. 2024. CAST Universal Design for Learning Guidelines version 3.0. <https://udlguidelines.cast.org/> Accessed: 2026-04-01.
- [12] Lijia Chen, Pingping Chen, and Zhijian Lin. 2020. Artificial Intelligence in Education: A Review. *IEEE Access* 8 (2020), 75264–75278. doi:10.1109/ACCESS.2020.2988510
- [13] Michelene TH Chi and Ruth Wylie. 2014. The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational psychologist* 49, 4 (2014), 219–243.
- [14] Thomas KF Chiu, Qi Xia, Xinyan Zhou, Ching Sing Chai, and Miaoqing Cheng. 2023. Systematic literature review on opportunities, challenges, and future research recommendations of artificial intelligence in education. *Computers and Education: Artificial Intelligence* 4 (2023), 100118.
- [15] Rudrajit Choudhuri, Dylan Liu, Igor Steinmacher, Marco Gerosa, and Anita Sarma. 2024. How Far Are We? The Triumphs and Trials of Generative AI in Learning Software Engineering. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering (Lisbon, Portugal) (ICSE '24)*. Association for Computing Machinery, New York, NY, USA, Article 184, 13 pages. doi:10.1145/3597503.3639201
- [16] Scott Crossley, Joon Suh Choi, Wesley Morris, Langdon Holmes, and David Joyner. 2025. AI Enhanced Intelligent Texts and Learning Gains. (2025).
- [17] Scott Crossley, Wesley Morris, Joon Suh Choi, and Langdon Holmes. 2025. Exploratory Assessment of Learning in an Intelligent Text Framework: iTELL RCT. In *Proceedings of the Twelfth ACM Conference on Learning@ Scale*. 2–12.
- [18] Rahul K Dass, Rochan H Madhusudhana, Erin C Deye, Shashank Verma, Timothy A Bydlon, Grace Brazil, and Ashok K Goel. 2025. Ivy: a hybrid knowledge-based and generative AI coach for explaining procedural skills. In *International Conference on Artificial Intelligence in Education*. Springer, 233–246.
- [19] Rocio De la Torre, Bhakti S Onggo, Canan G Corlu, Maria Nogueira, and Angel A Juan. 2021. The role of simulation and serious games in teaching concepts on circular economy and sustainable energy. *Energies* 14, 4 (2021), 1138.
- [20] Edward L Deci and Richard M Ryan. 2000. The "what" and "why" of goal pursuits: Human needs and the self-determination of behavior. *Psychological inquiry* 11, 4 (2000), 227–268.
- [21] Edward L Deci and Richard M Ryan. 2013. *Intrinsic motivation and self-determination in human behavior*. Springer Science & Business Media.
- [22] Brayan Diaz and Miguel Nussbaum. 2024. Artificial intelligence for teaching and learning in schools: The need for pedagogical intelligence. *Computers & Education* 217 (2024), 105071.
- [23] Adam Kenneth Dubé and Run Wen. 2022. Identification and evaluation of technology trends in K-12 education from 2011 to 2021. *Education and information technologies* 27, 2 (2022), 1929–1958.
- [24] Andrea Fedele, Clara Punzi, Stefano Tramacere, et al. 2024. The ALTAI checklist as a tool to assess ethical and legal implications for a trustworthy AI development in education. *Computer Law & Security Review* 53 (2024), 105986.
- [25] Centre for Excellence in Universal Design. 2026. The 7 Principles of Universal Design. <https://universaldesign.ie/about-universal-design/the-7-principles> Accessed: 2026-04-01.
- [26] Aaron A Funa and Renz Alvin E Gabay. 2025. Policy guidelines and recommendations on AI use in teaching and learning: A meta-synthesis study. *Social Sciences & Humanities Open* 11 (2025), 101221.
- [27] D Randy Garrison, Terry Anderson, and Walter Archer. 1999. Critical inquiry in a text-based environment: Computer conferencing in higher education. *The internet and higher education* 2, 2-3 (1999), 87–105.
- [28] D Randy Garrison, Terry Anderson, and Walter Archer. 2010. The first decade of the community of inquiry framework: A retrospective. *The internet and higher education* 13, 1-2 (2010), 5–9.
- [29] D Randy Garrison and J Ben Arbaugh. 2007. Researching the community of inquiry framework: Review, issues, and future directions. *The Internet and higher education* 10, 3 (2007), 157–172.
- [30] Ashok Goel, Chris Dede, Myk Garn, and Chaohua Ou. 2024. AI-ALOE: AI for reskilling, upskilling, and workforce development. *Ai Magazine* 45, 1 (2024), 77–82.
- [31] A Goel and D Joyner. 2015. Impact of a creativity support tool on student learning about scientific discovery processes. In *Proceedings of the Sixth International Conference on Computational Creativity*, Vol. 284.
- [32] Ashok K Goel and Lalith Polepeddi. 2018. Jill Watson: A virtual teaching assistant for online education. In *Learning engineering for online education*. Routledge, 120–143.
- [33] Adit Gupta and Christopher MacLellan. 2025. Intelligent tutors beyond K-12: An observational study of adult learner engagement and academic impact. *arXiv preprint arXiv:2502.16613* (2025).
- [34] Adit Gupta, Jennifer Reddig, Tommaso Calo, Daniel Weitekamp, and Christopher J MacLellan. 2025. Beyond final answers: Evaluating large language models for math tutoring. In *International Conference on Artificial Intelligence in Education*. Springer, 323–337.
- [35] Adit Gupta, Momin Siddiqui, Glen Smith, Jenn Reddig, and Christopher MacLellan. 2024. Intelligent tutors for adult learners: An analysis of needs and challenges. *arXiv preprint arXiv:2412.04477* (2024).
- [36] Golnoush Haddadian, Hyunkyu Han, Jinho Kim, Mohamed Shameer Abdeen, and Min Kyu Kim. 2025. Exploring AI-Generated Expert Models: Instructor Interaction and Learner Perceptions in a Physics Class. In *Proceedings of the 19th International Conference of the Learning Sciences-ICLS 2025*, pp. 1684–1688. International Society of the Learning Sciences.
- [37] Karen Holtzblatt and Hugh Beyer. 1997. *Contextual design: defining customer-centered systems*. Elsevier.
- [38] Sandeep Kakar, Rhea Basappa, Ida Camacho, Christopher Griswold, Alex Houk, Christopher Leung, Mustafa Tekman, Patrick Westervelt, Qiaosi Wang, and Ashok K Goel. 2024. SAM: an AI actor for fostering social interactions in online classrooms. In *International Conference on Intelligent Tutoring Systems*. Springer, 149–161.
- [39] Sandeep Kakar, Pratyusha Maiti, Karan Taneja, Alekhya Nandula, Gina Nguyen, Aiden Zhao, Vrinda Nandan, and Ashok Goel. 2024. Jill Watson: scaling and deploying an AI conversational agent in online classrooms. In *International Conference on Intelligent Tutoring Systems*. Springer, 78–90.
- [40] Hassan Khosravi, Simon Buckingham Shum, Guanliang Chen, Cristina Conati, Yi-Shan Tsai, Judy Kay, Simon Knight, Roberto Martinez-Maldonado, Shazia Sadiq, and Dragan Gašević. 2022. Explainable Artificial Intelligence in education. *Computers and Education: Artificial Intelligence* 3 (2022), 100074. doi:10.1016/j.caeai.2022.100074
- [41] Jinho Kim, Yoojin Bae, Jonathan Stravelakis, and Min Kyu Kim. 2024. Investigating the Influence of AI-Augmented Summarization on Concept Learning, Summarization Skills, Argumentative Essay Writing, and Course Outcomes in Online Adult Education. In *Proceedings of the 18th International Conference of the Learning Sciences-ICLS 2024*, pp. 2309–2310. International Society of the Learning Sciences.
- [42] Min Kyu Kim, Jinho Kim, and Ali Heidari. 2024. Exploring the multi-dimensional human mind: Model-based and text-based approaches. *Assessing Writing* 61 (2024), 100878.
- [43] Royce Kimmons and Joshua M Rosenberg. 2022. Trends and topics in educational technology, 2022 edition. *TechTrends* 66, 2 (2022), 134–140.
- [44] Malcolm S. Knowles. 1973. *The Adult Learner: A Neglected Species*. Gulf Publishing Company, Houston.
- [45] Malcolm Shepherd Knowles. 1984. *Andragogy in action*. (1984).
- [46] John Kos, Dinesh Ayyappan, and Ashok Goel. 2024. A Constructivist Framing of Wheel Spinning: Identifying Unproductive Behaviors with Sequence Analysis. In *International Conference on Intelligent Tutoring Systems*. Springer, 174–187.
- [47] J. Kos, R. Singh, C. Lum, and A. Goel. 2025. SCAFFOLDING LEARNERS' ECOLOGICAL DOMAIN KNOWLEDGE TO SUPPORT INQUIRY BASED MODELING. In *EDULEARN25 Proceedings (Palma, Spain) (17th International Conference on Education and New Learning Technologies)*. IATED, 5591–5599. doi:10.21125/edulearn.2025.1382
- [48] Lasha Labadze, Maya Grigolia, and Lela Machaidze. 2023. Role of AI chatbots in education: systematic literature review. *International journal of Educational Technology in Higher education* 20, 1 (2023), 56.
- [49] Robert Lindgren, Sandeep Kakar, Pratyusha Maiti, Karan Taneja, and Ashok Goel. 2024. Does Jill Watson Increase Teaching Presence?. In *Proceedings of the Eleventh ACM Conference on Learning@ Scale*. 269–273.
- [50] Marsha C Lovett. 1998. Cognitive task analysis in service of intelligent tutoring system design: A case study in statistics. In *International Conference on Intelligent Tutoring Systems*. Springer, 234–243.

- [51] Cherie Lum, Erin Deye, Grace Brazil, Tim Bydlon, Shashank Verma, Rochan Madhusudhana, Rahul Dass, and Ashok Goel. 2025. Designing an AI coaching system for interactive video-based skill learning. In *International Conference on Intelligent Tutoring Systems*. Springer, 281–291.
- [52] Jihao Luo, Chenxu Zheng, Jiamin Yin, and Hock Hai Teo. 2025. Design and assessment of AI-based learning tools in higher education: A systematic review. *International Journal of Educational Technology in Higher Education* 22, 1 (2025), 42.
- [53] Pratyusha Maiti and Ashok Goel. 2025. Can an AI Partner Empower Learners to Ask Critical Questions?. In *Proceedings of the 30th International Conference on Intelligent User Interfaces*. 314–324.
- [54] Nikola Marangunic and Andrina Granic. 2015. Technology acceptance model: a literature review from 1986 to 2013. *Universal access in the information society* 14, 1 (2015), 81–95.
- [55] Richard E Mayer. 1997. Multimedia learning: Are we asking the right questions? *Educational psychologist* 32, 1 (1997), 1–19.
- [56] Anja Meierkord and Roland Tusz. 2025. Trends in Adult Learning: New Data from the 2023 Survey of Adult Skills. Getting Skills Right. *OECD Publishing* (2025).
- [57] Thomas Mejtoft, Sarah Hale, and Ulrik Söderström. 2019. Design friction. In *Proceedings of the 31st European Conference on Cognitive Ergonomics*. 41–44.
- [58] Sharan B Merriam et al. 2001. Andragogy and self-directed learning: Pillars of adult learning theory. *New directions for adult and continuing education* 2001, 89 (2001), 3.
- [59] Robert J Mills and Matthew Harris. 2019. Alignment between technology acceptance and instructional design via self-efficacy. *Review of Business Information Systems* 23, 1 (2019).
- [60] Thomas Howard Morris. 2020. Experiential learning—a systematic review and revision of Kolb’s model. *Interactive learning environments* 28, 8 (2020), 1064–1077.
- [61] Wesley Morris, Joon Suh Choi, Langdon Holmes, Vaibhav Gupta, and Scott Crossley. 2024. Automatic Question Generation and Constructed Response Scoring in Intelligent Texts. In *Proceedings of the 17th International Conference on Educational Data Mining*.
- [62] Wesley Morris, Scott Crossley, Langdon Holmes, Chaohua Ou, Mihai Dascalu, and Danielle McNamara. 2025. Formative feedback on student-authored summaries in intelligent textbooks using large language models. *International Journal of Artificial Intelligence in Education* 35, 3 (2025), 1022–1043.
- [63] Wesley Morris, Scott Crossley, Langdon Holmes, Chaohua Ou, Danielle McNamara, and Mihai Dascalu. 2023. Using large language models to provide formative feedback in intelligent textbooks. In *International Conference on Artificial Intelligence in Education*. Springer, 484–489.
- [64] Chantal Mutimukwe, Olga Viberg, Lena-Maria Oberg, and Teresa Cerratto-Pargman. 2022. Students’ privacy concerns in learning analytics: Model development. *British Journal of Educational Technology* 53, 4 (2022), 932–951.
- [65] Andy Nguyen, Anh Thi Duong, Diep Thi Bich Nguyen, Van Thi Thanh Lai, and Belle Dang. 2025. Guidelines for learning design and assessment for generative artificial intelligence-integrated education: a unified view. *Information and Learning Sciences* (2025).
- [66] Andy Nguyen, Shohil Kishore, Yvonne Hong, Saima Qutab, and Belle Dang. 2025. Generative Artificial Intelligence (AI) in education: from organizing visions to official guidelines. *Information Technology & People* 38, 8 (2025), 172–199.
- [67] Jakob Nielsen and Rolf Molich. 1990. Heuristic evaluation of user interfaces. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 249–256.
- [68] EG Nikiforova and D Sh Shakirova. 2025. Ethical Guidelines for the Use of Artificial Intelligence in Education. In *International Scientific Conference Digital Future: Science, Education, and Innovative Development of Socio-Economic Systems*. Springer, 605–612.
- [69] JASMINE Park and AMELIA Vance. 2021. College Students’ Attitudes Toward Data Privacy. *Journal of Information Privacy and Security* 17, 2 (2021), 1–16.
- [70] Jan L Plass and Shashank Pawar. 2020. Toward a taxonomy of adaptivity for learning. *Journal of Research on Technology in Education* 52, 3 (2020), 275–300.
- [71] Rolf Ploetzner. 2024. The effectiveness of enhanced interaction features in educational videos: a meta-analysis. *Interactive Learning Environments* 32, 5 (2024), 1597–1612.
- [72] Valerie J Shute. 2008. Focus on formative feedback. *Review of educational research* 78, 1 (2008), 153–189.
- [73] Alexander Skulmowski and Kate Man Xu. 2022. Understanding cognitive load in digital and online learning: A new perspective on extraneous cognitive load. *Educational psychology review* 34, 1 (2022), 171–196.
- [74] Sharon Slade, Paul Prinsloo, and Mohammad Khalil. 2019. Learning analytics at the intersections of student trust, disclosure and benefit. In *Proceedings of the 9th international conference on learning analytics & knowledge*. 235–244.
- [75] Glen Smith, Adit Gupta, and Christopher MacLellan. 2024. Apprentice tutor builder: A platform for users to create and personalize intelligent tutors. *arXiv preprint arXiv:2404.07883* (2024).
- [76] Sergey Sosnovsky, Peter Brusilovsky, and Andrew Lan. 2025. Intelligent textbooks. *International Journal of Artificial Intelligence in Education* 35, 3 (2025), 967–986.
- [77] Brendan Strahm, Colin M Gray, and Mihaela Vorvoreanu. 2018. Generating mobile application onboarding insights through minimalist instruction. In *Proceedings of the 2018 Designing Interactive Systems Conference*. 361–372.
- [78] Karan Taneja, Pratyusha Maiti, Sandeep Kakar, Pranav Gururupasad, Sanjeev Rao, and Ashok K Goel. 2024. Jill watson: A virtual teaching assistant powered by chatgpt. In *International Conference on Artificial Intelligence in Education*. Springer, 324–337.
- [79] Karin Tusting and David Barton. 2003. *Models of adult learning: a literature review*. Niace UK.
- [80] Kurt VanLehn. 2011. The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational psychologist* 46, 4 (2011), 197–221.
- [81] Qiaosi Wang, Ida Camacho, Shan Jing, and Ashok K Goel. 2022. Understanding the design space of AI-mediated social interaction in online learning: challenges and opportunities. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW1 (2022), 1–26.
- [82] Qiaosi Wang, Shan Jing, Ida Camacho, David Joyner, and Ashok Goel. 2020. Jill Watson SA: Design and evaluation of a virtual agent to build communities among online learners. In *Extended abstracts of the 2020 CHI conference on human factors in computing systems*. 1–8.
- [83] Jennifer C Weeks and Lynn Hasher. 2014. The disruptive—and beneficial—effects of distraction on older adults’ cognitive performance. *Frontiers in psychology* 5 (2014), 133.
- [84] David Wood, Jerome S Bruner, and Gail Ross. 1976. The role of tutoring in problem solving. *Journal of child psychology and psychiatry* 17, 2 (1976), 89–100.
- [85] Weiqi Xu and Fan Ouyang. 2022. A systematic review of AI role in the educational system based on a proposed conceptual framework. *Education and Information Technologies* 27, 3 (2022), 4195–4223.
- [86] Hongxin Yan, Fuhua Lin, and Kinshuk. 2021. Including learning analytics in the loop of self-paced online course learning design. *International Journal of Artificial Intelligence in Education* 31, 4 (2021), 878–895.
- [87] Nur Yildirim, Changhoon Oh, Deniz Sayar, Kayla Brand, Supriya Challa, Violet Turri, Nina Crosby Walton, Anna Elise Wong, Jodi Forlizzi, James McCann, et al. 2023. Creating design resources to scaffold the ideation of AI concepts. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 2326–2346.
- [88] Ling Zhang, Richard Allen Carter Jr, Jeffrey A Greene, and Matthew L Bernacki. 2024. Unraveling challenges with the implementation of universal design for learning: A systematic literature review. *Educational Psychology Review* 36, 1 (2024), 35.

Received 19 January 2026; revised 12 March 2026; accepted 5 June 2026