

Reimagining Education in the Age of Generative AI

RESEARCH PAPER

Koan Research Team · September 14, 2025

ABSTRACT

The rapid rise of generative AI tools such as ChatGPT, Claude, and Microsoft Copilot has fundamentally reshaped education, transforming AI from a supplemental aid into a deeply embedded part of how students research, write, and process information. Research shows that heavy reliance on generative large language models can significantly diminish independent reasoning, writing, and problem-solving skills (Maes et al., 2025), a problem exacerbated by traditional educational structures that prioritize polished outputs over cognitive development. Koan addresses this challenge by integrating Aidan, a Socratic AI teaching assistant, into a learning management system. Aidan guides students through reflective problem-solving rather than providing answers, facilitating critical thinking and metacognition while giving educators insight into the learning process. By targeting cognitive, pedagogical, and institutional gaps, Koan aims to foster independent thought, meaningful engagement with AI, and a future-ready generation equipped for both academic and professional success.

Table of Contents

1. Introduction

2. Theory & Background

2.1 Cognitive

2.2 Pedagogic

2.3 Institutional & Future Consequences

3. Problem Statement

4. Our Product: Socratic AI for Public Cognitive Development

5. Why an AI-Integrated LMS?

6. Supporting Data: Internal Surveys

6.1 Student Findings

6.2 Educator Findings

6.3 Challenges, Concerns, and Desired AI Capabilities

7. Why Koan?

8. Conclusion

9. References

1. Introduction

The rapid emergence of generative artificial intelligence (AI) tools such as ChatGPT, Claude, and Microsoft Copilot has sparked a seismic shift in education. In the UK, 92% of students across multiple institutions say they are currently using some form of AI, up from 66% in 2024 (Freeman & Higher Education Policy Institute, 2025). Once viewed as supplementary, these tools are now embedded in how students research, write, and process information.

Notably, research shows that the manner with which students interact with generative LLM models is proving to be problematic. A 2025 study found that over a four-month period, “LLM users consistently underperformed at neural, linguistic, and behavioral levels” (Maes et al., 2025).

In many classrooms, this problem has been met not with integration, but with restriction. Schools and universities have responded with outright bans, and increased reliance on plagiarism detectors such as Turnitin and Unicheck, all inadequate attempts to contain what has already become a core part of students’ academic lives.

This institutional resistance misses the point. Our current education systems, having remained largely unchanged in their prioritization of output over critical learning, are simply not equipped to support learners in the burgeoning digital age. This represents the deeper structural problem; current educational systems have a damaged incentive structure. The product of polished assignments has become more important than the practice of learning, and students’ critical thinking skills are suffering for it.

The issue is not that students are using AI tools, but rather how they are using them. Therefore, our response must go beyond detection and discipline. It must address the cognitive, pedagogical, and institutional gaps that made generative AI such a common feature of student learning in the first place.

2. Theory & Background

Cognitive science, pedagogical theory, and institutional analysis each offer critical lenses for understanding how generative AI is reshaping education and what must change in response.

2.1 Cognitive

Cognitive science offers a clear framework for understanding the specific implications of generative AI in learning. Cognitive load theory (Sweller, 1988) tells us that learners can only handle a limited amount of information at once, so learning is most effective when distractions, or “extraneous load,” are eliminated. AI tools, when used appropriately, can alleviate this extraneous load by summarizing texts, organizing notes, or assisting with language translation. In this way, they are able to function as legitimate cognitive aids.

However, when AI is used to bypass “germane load,” i.e., the cognitive effort required to process, reflect on, and internalize new knowledge, the risk of surface learning increases. In the process of using LLMs to complete academic tasks, students produce acceptable outputs without engaging in the cognitive processes necessary for long-term retention or skill development. For example, when students rely on AI to generate a first draft of an essay, they skip the productive struggle of structuring their own arguments, which is a core component of critical thinking and knowledge construction.

Metacognition, or “the awareness and regulation of one’s own thinking,” is central to how students learn (Flavell, 1979). Students with low metacognitive awareness are more likely to rely on shortcuts or avoid cognitively demanding tasks (Zimmerman, 2002). In the context of generative AI, this means students may uncritically accept outputs, outsource reasoning, or skip reflection altogether. Without guided frameworks for ethical and reflective AI use, they risk reinforcing cognitive dependency and weakening the very skills (i.e., critical thinking, problem-solving, and self-regulated learning) that education is meant to cultivate.

2.2 Pedagogic

On a structural level, much of the current tension surrounding AI stems from how education systems define success. Standardized curricula, high-stakes testing, and grade-driven assessment models promote a form of learning that prioritizes outputs over process, incentivizing students to use generative tools in a way that is adverse to their long-term cognitive development (Au, 2007; Alfie Kohn, 2004; Shepard, 2000).

Digital pedagogy scholars present compelling data on why bans and restrictions are not the answer. They argue that the emphasis on surveillance and punishment in response to AI fails to recognize students as collaborators in their own learning (Morris & Stommel, 2018). Instead, they advocate for trust-based, process-centered learning environments that treat AI as a tool for dialogue and exploration rather than a threat to be policed. This is where Koan’s ethos aligns: rather than framing AI as a shortcut or a danger, our approach positions it as a partner in cognitive development by asking questions, prompting reflection, and supporting learning processes instead of replacing them.

Another important foundation is the Community of Inquiry (CoI) framework (Garrison, Anderson, & Archer, 2000). Developed at the beginning of the 21st century during a period marked by the rapid rise of computer-supported collaborative learning (CSCL) (Stahl et al., 2006), the framework proves timely, as we apply it to the similarly transformative emergence of AI in education. Like CSCL, CoI provides a bridge between cognitive theory and pedagogic design, offering a structured way to integrate new technologies while prioritizing meaningful learning.

The CoI framework highlights three dimensions of meaningful learning: (1) cognitive presence, the ability to make sense of new ideas through questioning and reflection; (2) social presence, the capacity for students to express themselves authentically and connect with peers; and (3) teaching presence, which involves designing and guiding the overall learning experience.

When applied to the integration of AI in education, intentionally and carefully designed platforms like Koan have the ability to strengthen all three presences. Koan can stimulate deeper cognitive presence by prompting students with structured, context-aware questions and follow-up probes that target the CoI phases of triggering event, exploration, integration, and resolution. Scaffolds such as “ask for evidence,” “compare alternatives,” “generate a counterexample,” and “plan next steps” encourage self-explanation and metacognition instead of giving pre-written answers. Koan can strengthen social presence by creating a space where students feel safe to “learn out loud.” They can make mistakes, test ideas, and engage in judgment-free dialogue that purposefully values genuine curiosity over polished performance. And finally, Koan can enhance teaching presence by providing controls that let instructors set how direct the Socratic tutor is. Example settings include hint frequency, depth of follow-up, when to reveal examples, and which rubrics or Bloom levels to target. Koan also includes dashboards that surface participation rates, distribution of question types, time on task, and indicators of the CoI phases. Prompt libraries tied to course outcomes and an alignment check that flags prompts or tutor moves that do not match the stated objectives further support this.

2.3 Institutional & Future Consequences

As of 2025, only 31% of American public schools have written policies on students' AI use (National Center for Education Statistics, 2025). Moreover, a whopping 82% of students reported using AI in their studies on a regular basis with less than half being able to say they had enough knowledge or skill to use generative AI responsibly (Dean, 2025). These statistics represent how educational institutions are being forced to navigate students' rapid adoption of AI technologies with inconsistent and slow-paced development of policies to manage their use in academic settings.

This disconnect undermines learning outcomes and erodes institutional and organizational confidence in students' capabilities. Public education exists to prepare citizens for economic participation, democratic engagement, and social contribution. When graduates lack critical thinking or the ability to work effectively with emerging technologies, these systems fail in their fundamental mission.

Without clear and effective AI policies, schools risk failing to prepare students for a technology-driven world. The current approaches, often inconsistent and restrictive, leave learners without guidance on ethical use, digital literacy, or productive collaboration with AI. Rather than equipping students with the skills needed to navigate emerging technologies, education systems are creating conditions where dependence on AI shortcuts replaces genuine learning, leaving graduates underprepared for both professional and civic challenges.

3. Problem Statement

“Modern education incentivizes polished outputs, not the messy, nonlinear process of learning. This has created a widening skills gap that erodes students' critical thinking and leaves students unprepared for real-world problem solving.”

4. Our Product: Socratic AI for Public Cognitive Development

At the heart of our public service solution lies Aidan, a Socratic AI teaching assistant built into an AI-integrated Learning Management System (LMS). Unlike conventional generative AI, Aidan is designed not to provide direct answers, but to guide students through questions and reflective learning processes. It guides them with questions and a reflective process that turns confusion into understanding. If a student asks for an essay on inflation, Aidan does not supply a draft. It breaks the task into decisions and prompts: What have you noticed about prices in your community this year? Which forces could explain those changes? How might public policy shape what families can afford? How would you organize an introduction, supporting paragraphs, and a conclusion so your claim is clear?

Under the hood, Aidan uses a Dynamic Knowledge Tree, a pedagogical frame that adapts to each learner while holding a consistent academic standard. Every exchange helps Aidan model how a student thinks, then ask the next best question. Process checkpoints, targeted hints, and formative feedback loops keep the student moving while preserving academic integrity. Teachers stay in control through ethical guardrails and adjustable guidance

levels, so they can set how directive or open the tutor should be. Early research shows that general-purpose AI used in isolation can harm learning. When placed in a safeguarded human-plus-AI workflow with clear teaching presence, it can increase engagement, comprehension, and critical thinking (Kestin et al., 2025; Bastani et al., 2024; Shukla & Pandey, 2025).

The Koan Platform

Koan turns that workflow into a single place to learn and produce work. Students draft, cite, revise, and submit inside one writing workspace. The editor supports rich text, equations, figures, citations, and version history. Aidan is available in the margin to outline, probe with questions, check rubric alignment, and turn feedback into a concrete revision plan. Submission uses snapshots, so the graded version stays fixed while students continue improving their live document.

Interactive assignments in Koan are structured activity forms that require real thinking and visible work inside the platform. Instead of upload-and-submit, each assignment is completed through steps inside Koan such as flowchart builders, Desmos-style graph tasks, interactive simulations, data tables to analyze, short rationale fields attached to each action, and reflection checkpoints. Inputs are atomic and time-stamped, with version history and event trails, which discourages outside AI use and makes learning traceable. Aidan sits beside each step to prompt reasoning without giving solutions, for example asking for a prediction before a simulation runs or an explanation after a graph is adjusted. Rubrics score process evidence as well as outcomes, and submission captures a snapshot of both the final work and the path taken. Assignments can also include peer review, paired roles, oral responses, and code or math widgets, all within the same guided flow.

Because Koan is an all-in-one platform, everything connects. The My Work hub shows what to do next across courses with one-click actions to continue writing or view feedback. Search understands courses, documents, comments, and assignments. Calendar aggregates deadlines and office hours and lets students start work from a date cell. Notifications and messages sit beside the work, not in a separate tool.

Teachers gain real-time visibility without extra overhead. Dashboards surface discourse quality, progress toward outcomes, and where students stall. Controls set the scope of allowed sources, the depth of hints, and when solutions can be revealed. Comment banks and quick rubrics speed high-quality feedback. Integrity tools show similarity views, citation checks, and reflection compliance.

In short, Aidan provides the questioning mind, interactive assignments provide the structure, and Koan provides the single place to do the work from first idea to graded submission to revision.

5. Why an AI-Integrated LMS?

As of 2024, over 270 million students, teachers, and institutional users globally utilize Learning Management Systems (LMS), with platforms like Canvas, Moodle, and Google Classroom leading the way. These LMS platforms, while widely adopted, do not natively support AI integration; any AI features are typically limited to assisting with navigation or administrative efficiency rather than enhancing the learning experience.

Existing commercial AI teaching assistants, such as Khan Academy's Khanmigo and MagicSchool AI, demonstrate the power of increased AI-human collaboration in learning and are examples of scalable success. However, they operate outside the LMS ecosystem, leaving teachers to juggle multiple tools and fractured workflows, thus limiting their effectiveness.

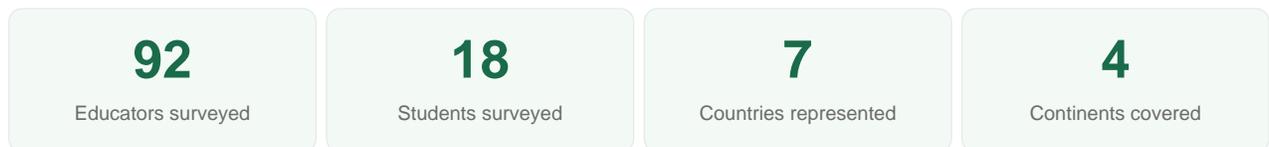
Herein lies a gap. Koan exists to fill this gap by integrating AI directly into the LMS, combining the benefits of both and responding to the evolving educational landscape. It addresses the need for personalized learning, efficient administrative processes, and enhanced student engagement, making AI an indispensable tool in modern education.

Research on interactive learning also consistently shows that “sustained implementation over one semester or longer” produces the greatest impact, particularly when combining mechanics (points, badges), dynamics (competition, collaboration), and aesthetics (narrative, immersion) (Kakar et al., 2024; Georgia Institute of Technology, 2023). This, fundamentally, is what Koan does.

Furthermore, UNESCO’s recently published AI Competency Framework for Students (Miao et al., 2024) emphasizes that developing AI fluency will become a necessary part of student learning and core competencies. Educators thus face a choice: lag behind the curve and wait for inconsistent and ill-fitting policies, or adopt this new learning paradigm that enhances both critical cognitive skills and AI competency.

6. Supporting Data: Internal Surveys

During Koan’s development and market validation, we sought to supplement our theoretical understanding by investigating the challenges educators and students face in an increasingly AI-driven learning environment, as well as the potential impact of these technologies on teaching and learning outcomes. To do this, we conducted two structured surveys: one targeting high school, undergraduate, and graduate-level educators (n = 92) across 7 countries and 4 continents, and another targeting students (n = 18) from 4 countries. Both surveys included a mix of Likert-scale, multiple choice, and open-ended questions, administered online in English and Spanish.



6.1 Student Findings

n = 18 (12 undergraduates, 5 secondary students, 1 graduate)

Our student survey, though small (n = 18), supported our hypothesis from research and lived student experience: AI is now deeply embedded in learning.

88% of students reported using AI tools (e.g., ChatGPT, Claude, Gemini) in the past year (95% CI: 73%–100%). 100% reported at least some AI use for assignments; not a single student selected “never.” Over 50% reported weekly or daily use (n = 12). The most common uses were writing or rephrasing text (83%), understanding difficult concepts (67%), and generating full answers or assignments (50%).

When asked about restrictions, 61% of students said bans on AI are ineffective, and 100% agreed that schools should teach students how to use AI ethically and effectively.

“Peers are using AI and completing assignments very quickly. If I don’t use AI, I feel like I am lagging behind.”

Undergraduate student, Evanston, Illinois, USA

Additionally, high academic pressure emerged as a recurring theme in responses. One high school student from Karnataka, India, described AI tools as a “less judgmental personal tutor.”

Perhaps most notably, 89% of students expressed concern that AI may be eroding their cognitive skills, with 94% saying they would be interested in trying a platform like Koan that focuses on guided thinking, rather than outsourcing answers. Though more responses are needed to strengthen the statistical significance of these results, they provide pre-emptive insights which illustrate that students are eager for guidance to use AI responsibly rather than simply banning it, signaling a clear need for a structured solution.

6.2 Educator Findings

n = 92 (primarily high school teachers and undergraduate lecturers)

Educator responses also highlighted both the challenges and opportunities presented by generative AI in the classroom. While students reported near-universal engagement with AI tools in their coursework, educators described a slightly more moderated but still substantial use: 75% of educators reported using AI tools to aid in planning, grading, or communication, compared with the 88% of students who reported using AI to complete assignments. This highlights the experiential differences: students interact with AI more directly and frequently in their learning, while many educators primarily use AI as a support tool rather than a learning partner.

Another point of divergence lies in disclosure practices. 49% ($n = 41$) of educators believe that students typically disclose their AI use, whereas student survey responses indicate that 100% of students do not always disclose usage. While only three students responded explicitly to the disclosure question (15 blank responses), these data perhaps point towards a skewed perception of student transparency and misaligned expectations between educators and learners. More student responses are required to provide further support for this conclusion.

6.3 Challenges, Concerns, and Desired AI Capabilities

Educators identified several core challenges and expressed clear ideas about how AI could best support teaching. 95% of educators ($n = 80$) report encountering at least one of the following problems with student AI usage: plagiarism, over-reliance replacing learning effort, inaccurate content, equity gaps, bias or harmful content, and distraction. Over-reliance replacing learning effort was the most frequently cited concern, followed by plagiarism.

In terms of adjustments to student assessments following the spread of student AI usage, 69% of educators reported changing their student assessments, including adjustments such as incorporating more oral presentations, follow-up questions, and handwritten work. Despite this, 70% disagreed that their workload had increased due to problematic AI usage, suggesting that they have found ways to positively incorporate AI tools in their pedagogy, and see opportunities for AI to support teaching rather than exacerbate administrative burdens.

About half of educators (47 respondents, 53%) combine multiple LMS rather than using a single platform. Educators use a variety of LMS tools to manage courses and communicate with students, but still highlight additional needs being unmet. They highlighted interest in AI tools that can provide individualized student support, break down problem-solving steps, and aid in lesson planning. However, educators were split on the idea of AI tools being used to assist in student grading. Some were firmly against it (primarily when providing feedback or comments on written work such as essays), while others welcomed its help in decreasing the grading workload.

The general consensus is that teachers would welcome AI assistance in correction or grading of simple quizzes or activities. One aspect of AI tool capability that instructors were very consistent and unified on was that AI should not replace human judgment, creativity, or the relational aspects of teaching.

Both surveys also confirmed the global scale of the issue. Responses came from educators and students in Brazil, the United States, Canada, India, Ecuador, and Australia, showing that the challenges Koan addresses are not confined to a single region or educational system. This international engagement demonstrates that over-reliance on AI and the need for structured guidance is a widespread concern, affecting diverse contexts and learner populations.

Interest in Koan

85%

Educators interested (95% CI: 77.5%–92.5%)

94%

Students interested (95% CI: 82.7%–100%)

Importantly, 78% of educators reported having full or shared authority over classroom AI policies, highlighting that adoption of a platform like Koan is feasible and timely. With demonstrated interest from students and educators alike, and its ability to address both cognitive and pedagogical challenges, Koan is well-positioned to transform AI from a shortcut into a tool for genuine learning. This suggests that despite challenges and variability in technology access, a solution like Koan could be globally adopted and highly relevant across diverse classroom contexts. The combination of urgent need, global relevance, and readiness for adoption indicates the current impact and extraordinary potential impact that Koan has.

7. Why Koan?

Ultimately, Koan is built for students, by students. In addition to being personally driven and motivated to sustain the platform, we are a dedicated team of university students from four continents with expertise in AI, research, and product development who have committed to develop this platform.

Our team has already completed ideation and preliminary market validation, demonstrating both the feasibility and demand for Koan. We are currently working on building early-stage partnerships, with strong student and educator interest driving momentum. Our interdisciplinary team combines expertise in sociology, education, and technology, giving us the capacity to scale thoughtfully and responsibly. We plan to continue growing through strategic partnerships with universities and schools, pilot programs, and an eventual subscription-based model that ensures long-term financial viability.

8. Conclusion

AI is no longer a distant possibility in education. It is already here, reshaping how students learn and interact with knowledge. Its presence is inevitable and permanent, making adaptation far more impactful than restriction. Bans and surveillance-focused measures fail to address the underlying challenges; what is required is a transformation in pedagogy that integrates AI thoughtfully into learning processes.

Canvas's strong market presence illustrates both the promise and the limitations of current LMS platforms: while widely used, they still struggle with usability issues and pedagogical shortcomings that innovative solutions can overcome. Platforms like Koan demonstrate that success lies in balancing sophisticated functionality with intuitive design, leveraging AI not as a shortcut but as a tool to enhance critical thinking, metacognition, and student engagement.

Ultimately, AI itself is not the problem. As stated earlier, the true challenge is our outdated approach to learning; we overvalue polished outputs and undervalue the messy, reflective process that builds deep understanding.

Koan addresses this by fostering active, process-centered learning while cultivating essential AI fluency. It's time to reimagine education for the AI era, with Koan.

9. References

- Alfie Kohn. (2004). *What does it mean to be well educated? and more essays on standards, grading, and other follies*. Beacon Press.
- Au, W. (2007). High-Stakes Testing and Curricular Control: A Qualitative Metasynthesis. *Educational Researcher*, 36(5), 258–267.
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabak, Ö., & Mariman, R. (2024). Generative AI Can Harm Learning. *The Wharton School Research Paper*.
- Dean, B. (2025). Percentage of Students That Use AI (2025). *Programs.com*.
- Flavell, J. H. (1979). Metacognition and Cognitive Monitoring: A New Area of Cognitive-Developmental Inquiry. *American Psychologist*, 34(10), 906–911.
- Freeman, J., & Higher Education Policy Institute. (2025). Student Generative AI Survey 2025. *Higher Education Policy Institute*.
- Garrison, D. R., Anderson, T., & Archer, W. (2000). Critical Inquiry in a Text-Based Environment. *The Internet and Higher Education*, 2(2–3), 87–105.
- Georgia Institute of Technology. (2024). Jill Watson – Design Intelligence Lab. *Gatech.edu*.
- Kakar, S., et al. (2024). Jill Watson: Scaling and Deploying an AI Conversational Agent in Online Classrooms. *Proceedings of ITS2024*.
- Kestin, G., Miller, K., Klales, A., Milbourne, T., & Ponti, G. (2025). AI tutoring outperforms in-class active learning. *Scientific Reports*, 15(1), 1–10.
- Kosmyna, N., et al. (2025). Your Brain on ChatGPT: Accumulation of Cognitive Debt when Using an AI Assistant for Essay Writing Task. *ArXiv*.
- Miao, F., Shiohira, K., & UNESCO. (2024). *AI competency framework for students*. UNESCO Publishing.
- Morris, S., & Stommel, J. (2018). *An urgency of teachers: the work of critical digital pedagogy*. Hybrid Pedagogy.
- National Center for Education Statistics. (2025). Press Release: More than Half of Public School Leaders Say Cell Phones Hurt Academic Performance. *NCES*.
- Shepard, L. A. (2000). The Role of Assessment in a Learning Culture. *Educational Researcher*, 29(7), 4–14.
- Shukla, H., & Pandey, K. (2025). Human-AI Collaboration in Teaching and Learning. *International Journal of Science and Social Science Research*, 2(4).
- Stahl, G., Koschmann, T., & Suthers, D. D. (2005). Computer-Supported Collaborative Learning. *Cambridge University Press EBooks*, 409–426.
- Sweller, J. (1988). Cognitive Load during Problem Solving: Effects on Learning. *Cognitive Science*, 12(2), 257–285.
- Zimmerman, B. J. (2002). Becoming a self-regulated learner: An overview. *Theory into Practice*, 41(2), 64–70.