Asian Journal of Education and Training
Vol. 11, No. 4, 198-210, 2025
ISSN(E) 2519-5387
DOI: 10.20448/edu.v11i4.7926

Comparing traditional AI, agentic ai and agentic rag for dialogic online education

Vincent English[1,2]

[1]*Longford International College, Ireland.*
[2]*Cambridge University, United Kingdom.*
Email: vincent.english@longfordcollege.com

Abstract

Online education increasingly depends on artificial intelligence (AI) for scale, personalization, and assessment. However, most deployments remain confined to one-shot, content-delivery paradigms that under-serve dialogic pedagogy, an approach centered on multi-voiced inquiry, co-construction of knowledge, and iterative, socially mediated reasoning. This paper synthesizes three paradigms of AI: Traditional AI, Agentic AI, and Agentic Retrieval-Augmented Generation (RAG), and evaluates how each can be applied to online teaching and learning organized around dialogic principles. I articulate a theory-led design space grounded in dialogic pedagogy (Freire, Bakhtin, Wegerif, Alexander) and contemporary learning science (Vygotsky's ZPD; the Community of Inquiry framework; ICAP). I map each AI paradigm to core online education tasks (tutoring, assessment for learning, discussion orchestration, knowledge building). I propose reference architectures and governance patterns and offer implementation roadmaps, metrics, and risk mitigations. The paper argues that while traditional AI enables efficient, bounded tasks (e.g., automated grading, item generation), agentic AI introduces goal-directed orchestration across tools and actions required for authentic dialogic workflows (e.g., facilitation, critique, reflection). Agentic RAG best aligns with dialogic pedagogy by grounding agent decisions in evolving, cited knowledge; supporting multi-turn planning and verification; and maintaining memory of class discourse and norms. The paper concludes with a pragmatic recommendation: combine Agentic RAG for knowledge-intensive, discourse-heavy learning with narrowly scoped traditional AI services and agentic guards; evaluate with dialogic outcome metrics, not merely accuracy or time-on-task.

Keywords: Agentic AI, Community of Inquiry, Dialogic Pedagogy, Educational Governance, Formative Assessment, ICAP, Instructional Design, Learning Analytics, Online Education, Retrieval-Augmented Generation.

Citation | English, V. (2025). Comparing traditional AI, agentic ai and agentic rag for dialogic online education. *Asian Journal of Education and Training, 11*(4), 198–210. 10.20448/edu.v11i4.7926
History:
Received: 4 November 2025
Revised: 8 December 2025
Accepted: 15 December 2025
Published: 24 December 2025

Publisher: Asian Online Journal Publishing Group

Funding: This study received no specific financial support.
Institutional Review Board Statement: Not applicable.
Transparency: The author confirms that the manuscript is an honest, accurate, and transparent account of the study; that no vital features of the study have been omitted; and that any discrepancies from the study as planned have been explained. This study followed all ethical practices during writing.
Competing Interests: The author declares that there are no conflicts of interests regarding the publication of this paper.

Contents

Contribution of this paper to the literature
This paper introduces a dialogic pedagogy–based design framework that compares Traditional AI, Agentic AI, and Agentic RAG for online education. It maps each paradigm to core teaching tasks, specifies reference architectures and governance patterns, and proposes implementation roadmaps, evaluation metrics, and risk mitigation strategies. The paper recommends a pragmatic hybrid approach centered on Agentic RAG.

1. Introduction

The digital transformation of higher education, accelerated by the COVID-19 pandemic, has increased access to education but also highlighted the limitations of prevailing online pedagogical models, which often reproduce a transmission-style, content-delivery approach (Bozkurt et al., 2020). This stands in stark contrast to dialogic pedagogy, which views learning as a social, multi-voiced process of inquiry and co-construction of knowledge (Alexander, 2008; Bakhtin, 1981; Freire, 1970; Wegerif, 2007). Authentic dialogue at scale online requires systems that can do more than deliver content; they must be able to plan, listen, retrieve, and cite information, facilitate multiple perspectives, and adapt to class norms. Developments in agentic AI and retrieval-augmented systems seem to offer a promising pathway to meet these challenges. In this respect, this paper will examine three distinct AI paradigms.

- Traditional AI: This paradigm involves training a model on a specific dataset for a narrow, well-defined task, and then deploying it as an inference service. Examples include automated essay scoring, plagiarism detection, and generating multiple-choice questions from a text.
- Agentic AI: This paradigm moves beyond single-task execution. Here, large language models (LLMs) act as the "brain" of an agent that can reason, plan a sequence of actions, and utilize a variety of digital tools and APIs (e.g., search engines, calendars, LMS functions) to accomplish more complex, multi-step goals (Chu et al., 2025).
- Agentic Retrieval-Augmented Generation (RAG): This is a specialized form of Agentic AI that addresses the inherent limitations of large language models (LLMs), such as knowledge cut-offs and the tendency to "hallucinate" or invent information. An Agentic RAG system continuously retrieves information from a dynamic, trusted knowledge base, grounds its generated responses in that evidence, and explicitly cites its sources. Crucially, for dialogic learning, it can maintain a memory of interactions over time, allowing it to adapt to the specific context of a class (Li et al., 2025).

This paper provides a comparative analysis of these paradigms through the lens of dialogic pedagogy, addressing three key questions.

1. How do the capabilities and limitations of Traditional AI, Agentic AI, and Agentic RAG map onto the core processes of dialogic teaching and learning in online environments?
2. What reference architectures, governance policies, and evaluation metrics are necessary to support the trustworthy, equitable, and pedagogically effective deployment of these AI paradigms?
3. Which paradigm or combination of paradigms is most suitable for fostering a rich dialogic environment online, and under what specific constraints and conditions?

I will argue in this paper that Agentic RAG can offer a suitable scaffolding for the dialogic goals in education, provided that it is supported by robust academic governance, learning-centered formative analytics, and human-in-the-loop oversight. This paper will outline the theoretical foundations, present a conceptual framework, compare the three paradigms, and propose reference architectures, an implementation roadmap, and a discussion of ethical considerations.

2. Literature and Theoretical Background

This section combines key theoretical frameworks from dialogic pedagogy. It will explore Bakhtin's understanding of dialogism, then Freire's concept of an 'anti-dialogic' 'banking' approach to education, followed by Alexander's key aspects of dialogic education and conclude with a final examination of dialogic spaces from Wegerif.

2.1. The Foundations of Dialogic Pedagogy

Dialogic pedagogy holds that learning is a transformative process in which ideas are tested and refined within a learning space between interlocutors (Wegerif, 2007). There are four main contributors to the academic discussion on dialogism. Bakhtin posits that dialogism is polyphonic, in which every utterance responds to prior utterances and also anticipates future ones. Thus, meaning is created through the dynamic interplay of voices across different times and social contexts, which exist in the dialogic space thus created (Bakhtin, 1981; Bakhtin, 1986). Hence, within the dialogic space, a student's utterance never stands alone. Accordingly, a dialogic space is created in response to another interlocutor, such as a teacher to a student, or to the texts they have read. An AI designed for dialogic learning must, therefore, be able to situate and interpret contributions within this ongoing web of discourse.

Secondly, Freire (1970) criticized the transmissive concept of "banking" education. He defines this as teachers who 'transmit' knowledge to passive students, who have it 'deposited' in their minds. To counter this, Freire suggests a "problem-posing" approach where students become co-creators and co-investigators. Knowledge and learning, therefore, begin with generative themes defined from lived realities, which are then explored and examined through active dialogism. This process, which Freire called conscientização (critical consciousness), empowers learners to question their world and become agents of their own transformation. Alexander (2008) described five features of effective dialogic teaching. He viewed it as a process that moves from theory to classroom application. The first feature is collective talk, where students and teachers work together on tasks, followed by 'reciprocal talk'. Both students and teachers listen, share ideas, and consider each other's perspectives. The next aspect is supportive talk, where learners can speak freely without fear of embarrassment or criticism. Building on the first three elements is cumulative talk; here, interlocutors build on their own and each other's ideas. The final aspect is purposeful talk, in which teachers plan and guide discussions with clear objectives. These five features provide teachers with a straightforward framework to design and evaluate classroom dialogue. This applies whether a person or a computer leads the discussion.

Finally, Wegerif (2007) argues that the goal of education is not just to acquire knowledge but to expand one's "dialogic space" which he defines as a 'space' that scaffolds the capacity to engage with and learn from otherness and new possibilities. This necessarily involves a shift in identity, from seeing oneself as a container of knowledge to seeing oneself as a participant in an ongoing dialogue. Technology, in this view, should be designed not merely to transmit information but to open and hold this dialogic space. An AI for dialogic learning should therefore be designed to encourage perspective-taking, intellectual humility, and an orientation towards understanding rather than winning an argument. It should not be built merely to transmit information but to open and sustain a dialogic space in which inquiry can unfold. An AI designed for dialogic learning would cultivate perspective-taking by inviting learners to consider alternative viewpoints, empathize with others' positions, and interrogate their own assumptions. It would also model and prompt intellectual humility, normalizing the acknowledgment of uncertainty and the willingness to revise one's stance in light of better reasons or new evidence. Equally, such a system would orient interaction toward understanding rather than victory, structuring exchanges so that the goal is clarification, sense-making, and co-construction of meaning rather than rhetorical triumph. To make this possible, it would maintain a safe and inclusive environment in which all participants feel both welcomed and accountable—comfortable enough to risk ideas and challenge the status quo yet guided by norms of respect that keep disagreement productive. In combination, these design commitments position AI not as a conduit for content delivery but as an active steward of dialogic practice (Figure 1).

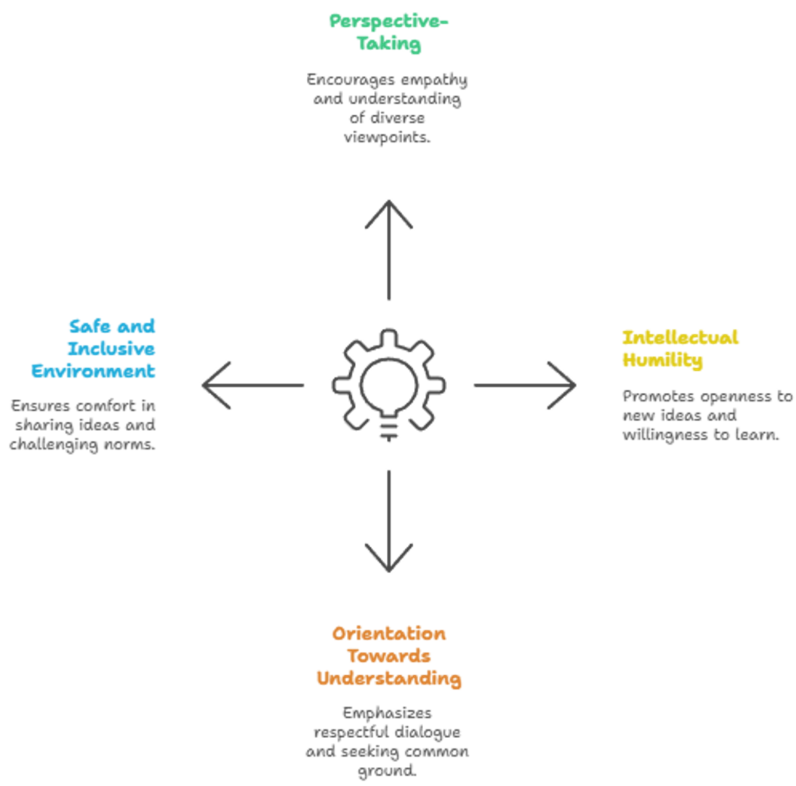

Figure 1. Cultivating open-mindedness.

Taken together, these perspectives establish a clear set of values for a dialogically oriented education. It prioritizes multi-turn exchange over single-shot Q&A, civil contestation, and the exploration of multiple viewpoints over consensus. Reasoning grounded in evidence is preferred over unsupported assertions, and the co-authoring of shared understanding and artifacts is valued over individual performance.

2.2. Learning Science Anchors for Online Settings

Dialogic pedagogy in online education is underpinned by a set of learning-science traditions that together specify how interaction should be structured, supported, and assessed. Vygotsky's account of the Zone of Proximal Development (ZPD) locates learning in the space between what a learner can do unaided and what becomes possible with guidance; scaffolding is the disciplined provision and gradual withdrawal of that support (Vygotsky, 1978). As mentioned earlier, AI can serve as a scaffolding partner in a dialogic space by offering hints, worked examples, and opportunities for inquiry. Teachers can use Socratic prompts tailored to the learner's evolving capabilities, with agentic systems adjusting support dynamically in response to real-time performance.

The well-established Community of Inquiry (CoI) framework complements this view by specifying the conditions under which meaningful learning emerges in online environments. These conditions include cognitive, social, and teaching presence that operate in concert to sustain inquiry and shared understanding (Garrison, Anderson, & Archer, 2000). Accordingly, AI can strengthen and support the 'teaching presence' by automating routine facilitation and curating resources, widen social presence by orchestrating timely peer interaction, and deepen cognitive presence by structuring dialogue around evidence and argumentation (Anderson, Nguyen, & Moreira, 2025). From the perspective of engagement quality, the ICAP framework further clarifies desired learner activity patterns, distinguishing passive and active behaviors from the more potent constructive and interactive modes. Dialogic pedagogy explicitly seeks to elicit the latter, for example, by transforming solitary reading into an AI-mediated exchange that encourages explanations, counterexamples, and co-construction of ideas (Chi & Wylie, 2014).

Assessment and design considerations complete the picture. Formative assessment research demonstrates that timely, specific, and actionable feedback is among the most powerful levers for improvement; the aim is not to grade but to guide the next steps for both the learner and the teacher (Black & Wiliam, 1998). Large language models can deliver rapid, fine-grained commentary on diverse artefacts from short answers to extended essays provided that feedback is framed dialogically to prompt reflection, revision, and transfer. At the same time, Cognitive Load Theory reminds us that working memory is finite and that online experiences must minimize extraneous load so that intrinsic and germane processing can proceed (Sweller, 1988). Personalized pathways, just-in-time cues, and the automation of peripheral tasks are practical ways AI can reduce unnecessary burdens while maintaining productive challenges. Taken together, these principles serve as both constraints and goals for the design of AI in dialogic online education: systems should scaffold learners within their Zone of Proximal Development (ZPD), cultivate a robust community of inquiry with clear teaching presence, advance activities up the ICAP hierarchy toward interactive knowledge building, embed feedback that is genuinely formative and conversational, and manage cognitive load so that attention is reserved for the work of understanding.

2.3. The Evolution of AI Paradigms in Education

AI in education has progressed through distinct phases, moving from narrow statistical models optimized for specific, well-defined tasks to today's large language models (LLMs) capable of open-ended dialogue, reasoning, and content generation. Early "Traditional AI" systems (e.g., rule-based engines and supervised models) excelled at bounded tasks such as automated data scoring, item generation, and basic analytics. The problem was that they operated primarily as black-box utilities, detached from the unfolding pedagogy, with nothing visible in a pedagogical sense. The emergence of LLMs changed this paradigm, enabling users to engage in conversations and, in turn, receive formative AI tutoring and feedback. On the debit side, AI used naively risks ungrounded claims (hallucination) and brittle behavior outside their training data. In response, "Agentic AI" adds planning and control: it can set goals, chain tools, monitor progress, and escalate to humans, thereby coordinating the processes that make online learning social and iterative building further, "Agentic RAG" (Retrieval-augmented generation) couples these orchestration strengths with explicit grounding in curated and evolving corpora, producing cited, auditable interventions that align more closely with dialogic, evidence-based education. In what follows, I consider each paradigm in turn, its characteristic capabilities, limitations, and design implications for trustworthy, equitable, and pedagogically effective online learning.

2.3.1. Traditional AI

This initial wave of educational artificial intelligence focused on well-defined, repetitive tasks. Systems were built using machine learning models trained for specific purposes, such as Intelligent Tutoring Systems (ITS) that guide students through structured problem sets. These automated essay scoring systems predict a human grader's score based on textual features, and recommendation engines that suggest learning resources based on a user's profile. These systems excel at providing scalable, consistent solutions for stable tasks. However, they are brittle; they cannot handle tasks outside their training distribution and cannot engage in the kind of open-ended, multi-turn dialogue that is characteristic of dialogic pedagogy.

2.3.2. Agentic AI

The advent of powerful large language models (LLMs) like those in the GPT series has led to the emergence of agentic AIs. Agentic agents are specialized systems capable of perceiving their specific environment and reasoning about how to achieve a particular goal. Essentially, this allows specialized LLMs to focus on specific tasks, such as mathematics tutoring (Chu et al., 2025). In this context, an LLM functions as the agent's "brain," formulating a plan to accomplish a user-defined objective, for example, organizing a week-long online debate on climate change for a class. The ability to coordinate a sequence of actions across multiple tools, such as using APIs, marks a significant advancement in capability. This development enables the automation of complex pedagogical workflows that were previously the exclusive domain of human teachers, which was not possible before the advent of LLMs.

2.3.3. Agentic Retrieval-Augmented Generation (RAG)

Agentic AI may well be powerful, but it inherits the fundamental weaknesses of the large language models (LLMs) that power it because LLMs are trained on static datasets, meaning their knowledge is outdated until the next training. In the meantime, LLMs "hallucinate," producing output that is inaccurate or entirely false. Agentic Retrieval-Augmented Generation (RAG) is therefore designed to mitigate these risks (Li et al., 2025). Before generating a response, the agent retrieves relevant information from an authoritative, up-to-date knowledge source (e.g., a vector database containing course readings, recent academic papers, or news articles). The retrieved information is then passed to the LLM as part of the prompt, instructing it to use this information to formulate its response and to cite its sources. This process of grounding the generation in retrieved evidence dramatically reduces hallucinations and increases the trustworthiness of the output. Furthermore, the system can be designed to include a memory that stores the history of interactions (e.g., previous questions, student misconceptions, class norms) in its knowledge base. This allows the agent to adapt its behavior over time, making it a much more suitable partner for the ongoing, cumulative nature of dialogic knowledge building. This evolution from static, single-task AI to dynamic, grounded, and memory-enabled agentic systems opens up new possibilities for supporting dialogic pedagogy online. The remainder of this paper will explore these possibilities in detail. I now comparatively analyze how the three AI paradigms apply to the core tasks of dialogic online education, beginning with a high-level comparison and then examining specific design patterns.

2.4.High-Level Paradigm-Task Fit

To clarify how different technical approaches translate into classroom practice, Table 1 contrasts three broad AI paradigms: Traditional AI, Agentic AI, and Agentic RAG against the main work of dialogic online education. It summarizes, at a high level, how each paradigm supports key tasks such as initiating inquiry, steering discussion, tutoring, knowledge building, assessment, inclusion, and governance. The table is not intended as a definitive ranking but as a heuristic overview of where each paradigm is currently most (and least) aligned with the demands of dialogic teaching and learning.

Table 1. Paradigm–Task fit in dialogic online education.

Task / Capability	Traditional AI	Agentic AI	Agentic RAG
Launching inquiries & provocations	Mostly limited to serving up fixed prompts, question banks, or simple quiz stems drawn from pre-authored materials.	Can design and run short sequences of prompts or activities, schedule when they appear, and organize students into groups using basic profile information.	Uses up-to-date, cited material from curated collections and the web to seed inquiries, so that starting questions are timely, contested, and anchored in real-world debates.
Orchestrating discussion	Relies on simple, pre-set rules (e.g., keyword flags or turn-taking scripts) and struggles to respond to the evolving shape of a conversation.	Actively tracks the flow of discussion, encourages quieter students, assigns rotating roles (e.g., summarizer, challenger), and alerts a human instructor when norms or community guidelines appear to be at risk.	Builds on the orchestration abilities of Agentic AI but can also connect contributions to shared class resources and external texts (e.g., "How does this comment relate to last week's article?"), making prompts more substantively grounded.
Tutoring & formative feedback	Works well for item-level feedback, hints, and scoring structured responses against a rubric, especially for patterns it has been explicitly trained on.	Devises multi-step improvement plans (e.g., review, practice, reflection), routes learners to suitable resources, and determines when a human tutor should intervene.	Offers dialogic feedback that references specific sources, takes into account the learner's prior misconceptions or attempts, and checks its own explanations against a knowledge base before presenting them.
Knowledge building & synthesis	Can generate short summaries of threads or documents but often misses nuance and struggles to integrate multiple perspectives or sources.	Coordinates the process of synthesis by assigning tasks (e.g., different students summarise different viewpoints) and then combining these contributions into a provisional overview.	Maintains a living, referenced knowledge artifact (e.g., a shared document or concept map) that traces claims, counterclaims, and gaps over time, helping the class see emerging consensus, tensions, and open questions.
Assessment for Learning	Provides reliable automated scoring for closed or highly structured tasks and basic analytics on activity levels or completion.	Draws on evidence from varied sources (discussion posts, drafts, logs) to build a richer picture of learner progress, with an emphasis on process and growth.	Adds checks on epistemic quality for instance, examining citation practices, probing for counterevidence to a claim, and generating prompts that ask students to evaluate their own work against a rubric.
Inclusion & accessibility	Offers basic personalisation, such as adjusting difficulty or providing alternative question formats.	Chains together tools and workflows, such as translation and text-to-speech, to support students with specific accessibility or language needs.	Tailors explanations and examples using culturally and linguistically diverse sources, while considering individual preferences and barriers, to ensure that interactions remain responsive over time.
Governance & auditability	High transparency: deterministic models and tightly scoped tasks make decisions relatively straightforward to inspect and explain.	Autonomy increases complexity, but actions can still be reconstructed from logs of tool calls and decision steps.	Although it is more complex, the use of cited sources provides a clear trail to follow. You can see exactly where key claims originate, which makes it much easier to verify, question, and review the system's reasoning.

2.5. Summary of Findings

The analysis demonstrates a clear, step-by-step progression. Traditional AI is effective at managing small, repetitive tasks; however, it does not significantly facilitate collaborative communication or collective thinking. Agentic AI advances this by assuming parts of the dialogue process itself, effectively acting as a co-facilitator in the learning environment. Agentic Retrieval-Augmented Generation (RAG) most directly supports dialogic epistemics, building knowledge through claims, evidence, and synthesis by grounding interactions in a verifiable knowledge base and maintaining a community memory. To translate conceptual discussions into practical plans, I outline reference architectures for each AI paradigm, describing key technical components and their interactions.

2.5.1. Traditional AI in the Learning Management System (LMS)

The architecture for traditional AI is a service-oriented model where AI functionalities are exposed as discrete endpoints callable by the LMS. This represents a standard and straightforward integration model.

- Core Components: The core of this architecture consists of a set of deployed machine learning models, each trained for a specific task. These models are wrapped in APIs, creating endpoints for services such as score_essay, tag_discussion_post, and generate_quiz_questions. These services are designed to be stateless and deterministic.
- Data and Models: The models are trained on large, static datasets. For an essay scorer, this would be a corpus of essays previously graded by humans. For a quiz generator, it would be a collection of texts and their corresponding questions. The data used for training, such as item banks, rubrics, and anonymized student discussion transcripts, is crucial for the model's performance and fairness.
- Governance and Integration: Governance in this model focuses on the initial validation and periodic recalibration of the models. This involves regularly checking the training data for bias, sharing clear documentation (model cards) that explain how the model works, where it performs well, and where it does not, and implementing a human review step for important AI outputs. In practice, the system typically connects to the LMS through a simple REST API: the LMS sends the relevant data (for example, a student's essay), and the AI responds with a score and confidence rating in JSON format.

2.5.2. Agentic AI Orchestration

An agentic AI architecture interposes a stateful planning and execution layer between user intent and the underlying computational affordances. At its core is a planner (or agent), typically instantiated by an LLM, that translates high-level instructional goals such as "facilitate this week's discussion" into a decomposed sequence of executable actions, selecting appropriate tools for each step from a predefined library. This library encompasses not only traditional AI micro-services (e.g., scoring, item generation) but also application interfaces for the learning management system (e.g., creating forum posts, retrieving rosters), institutional platforms, and external utilities such as web search or messaging.

To operate coherently over time, the agent maintains a structured memory substrate. Short-term memory functions as a scratchpad for the current plan and intermediate results; session memory preserves the local context of an ongoing interaction (for example, the assignment a student is working on); and long-term memory records durable information about learners and the environment, supporting personalization, continuity, and cumulative pedagogy. Surrounding these components is an observation-and-control layer that ensures governance and safety: comprehensive action tracing, rate limiting of tool calls, and explicit human-in-the-loop checkpoints for high-stakes operations (such as broadcasting class-wide messages). Together, these elements yield a system capable of coordinating complex educational workflows while remaining auditable, policy-compliant, and pedagogically accountable.

2.5.3. Agentic RAG for Dialogic Pedagogy

At the core of an Agentic Retrieval-Augmented Generation (RAG) system is the retriever (Figure 2), which functions as an intelligent search engine capable of rapidly locating relevant information across various knowledge sources. The system typically employs vector search techniques, enabling it to match ideas based on their semantic meaning rather than relying solely on keyword matching. The retriever operates across three primary corpora. The first is the course corpus, encompassing all materials officially associated with the course, such as the syllabus, assigned readings, lecture notes and transcripts, slides, and any additional resources provided by the instructor. This corpus serves as the main reference for questions directly related to the course content. The second is the class corpus, a dynamic archive capturing real-time classroom interactions, including discussion posts, question-and-answer exchanges, chat logs, and collaborative activities. Over time, this corpus evolves into a comprehensive record of the group's collective thinking, misconceptions, and insights, continually expanding as new conversations occur. The third is a curated external corpus, composed of carefully selected, reputable outside sources such as open-access journal articles, credible news outlets, industry reports, and other authoritative materials. Incorporating this external corpus allows the system to incorporate broader perspectives and current examples, thereby extending the scope of inquiry beyond the official course materials.

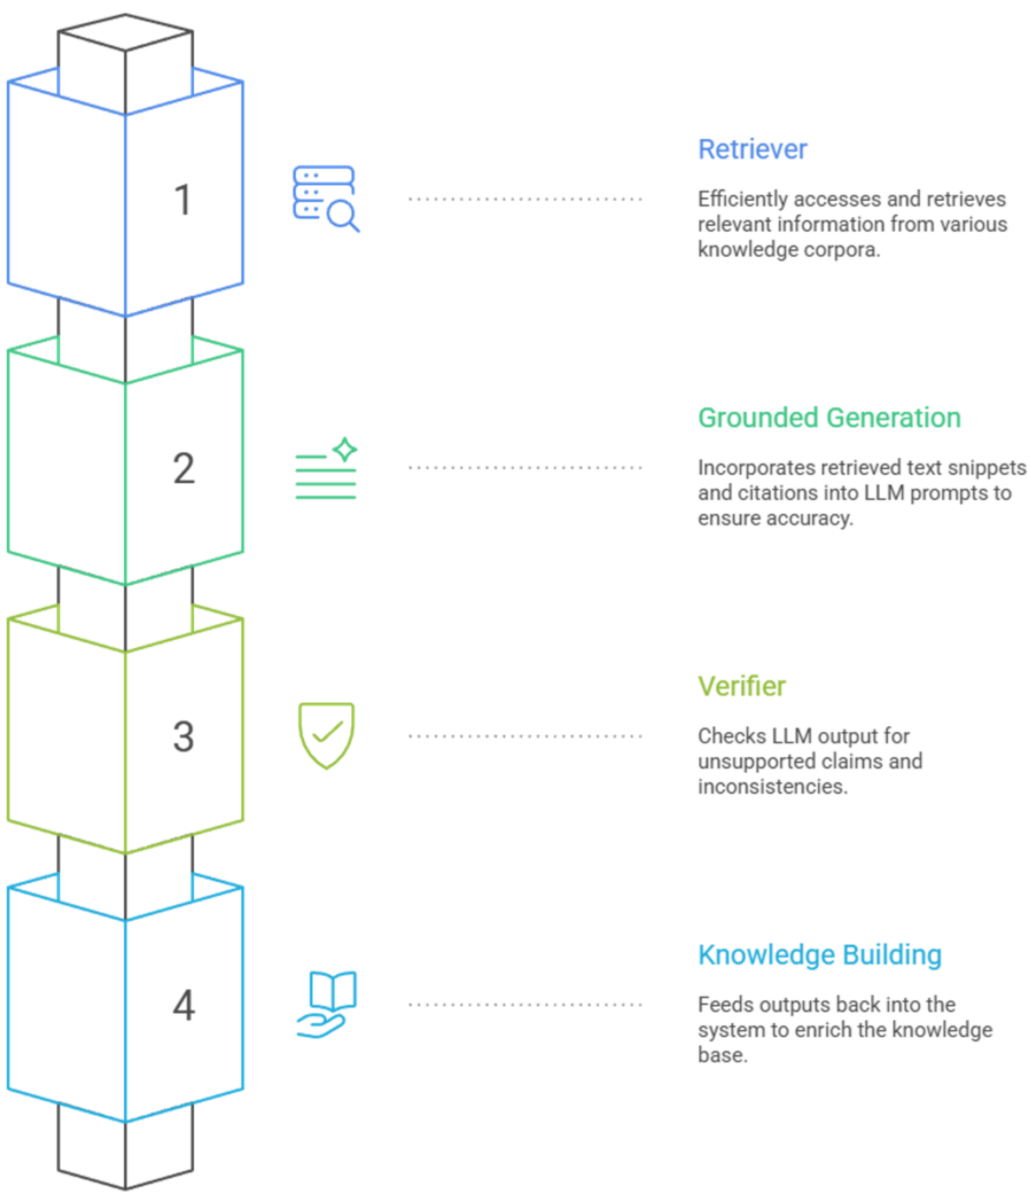

Figure 2. Building a robust Agentic RAG architecture.

Taken together, these architectures reflect a clear progression in both technical complexity and educational impact. Traditional AI is relatively straightforward to implement, whereas a fully developed Agentic RAG system requires significantly more investment in infrastructure, data, and design. However, it also offers the greatest potential for genuinely dialogic, knowledge-rich online learning. At the same time, integrating powerful AI systems into education raises serious ethical and governance questions. Transitioning from deterministic, rule-bound tools to more autonomous, agentic systems makes it essential to establish strong governance structures that emphasize safety, fairness, transparency, and meaningful human oversight. The following section outlines the main risks and proposes a multi-layered response to managing them. This is not merely a theoretical concern: many institutions are already deploying AI tools without clear governance, increasing the likelihood of inconsistent and inequitable use (RNL, 2024). In this context, an ethical framework is not optional but a necessary foundation for any sustainable and responsible innovation.

2.6. Risks and Mitigations

The risks associated with educational AI vary depending on the paradigm. The following table expands on the initial risk assessment by incorporating recent research (Chau, Yan, & Liu, 2025; Slimi & Carballido, 2023) provide insights into these risks. Table 2 synthesizes the distinctive risk landscape that emerges when Traditional AI, Agentic AI, and Agentic Retrieval-Augmented Generation (RAG) are deployed in dialogic education, pairing each risk with pragmatic safeguards. It highlights how core hazards such as hallucination and misinformation, bias and representation, privacy and data governance, over-automation and de-skilling, academic integrity and gaming, and explainability and auditability manifest differently across the three modalities. These range from bounded errors in conventional systems to tool-mediated exposure and autonomy risks in agentic systems, and corpus-dependent vulnerabilities in retrieval-augmented designs. Additionally, the table organizes mitigation strategies along policy, technical, and process dimensions. These include enforceable standards like citation and transparency requirements, verifiable controls such as auditors, verifiers, and access governance, and human-centered routines like teaching presence and assessment redesign. Taken together, the matrix serves as a design brief: it encourages leveraging the strengths of each approach while implementing layered protections that uphold dialogic values such as equity, safety, intellectual humility, and shared inquiry.

Table 2. Expanded risks and mitigations for AI in dialogic education.

Risk Category	Traditional AI	Agentic AI	Agentic RAG	General mitigations and policy recommendations
Hallucination & Misinformation	The overall risk is relatively low and mostly confined to the training data. Errors typically show up as bounded mistakes, such as a misclassified response or an inaccurate score on a specific task.	Risk increases as the agent starts calling external tools (e.g., web search) without firm grounding. This can introduce unchecked or misleading information into the learning environment.	The RAG setup is explicitly designed to reduce hallucination by grounding responses in retrieved sources, but the risk does not disappear. It resurfaces if the underlying corpus is of poor quality or the generator fails to stick closely to the retrieved context.	In practice, this means asking the AI to show its workings. Any factual claim it makes especially in high-stakes situations should be backed up with a clear citation. On the technical side, its answers should be checked by verifier systems or secondary checks that compare what it says against more than one source. At the process level, there should always be a human involved in important decisions: key pieces of feedback are reviewed by a person, and anything the AI is not very confident about is automatically flagged for human checking before it is used.
Bias & Unfair Representation	Bias risk is high and largely driven by the training data. If the data over-represents certain groups or viewpoints, model performance may be less accurate—and potentially unfair—for others.	Bias can appear not only in data but also in how the agent chooses tools and plans its actions. For example, it might favour certain sources or patterns of support, unintentionally privileging particular student profiles.	Here, the biggest issue becomes how we choose and maintain the materials the system can search. Even if each piece of information is accurate, if the corpus leans too heavily toward certain viewpoints or cultures, the system will keep mirroring and amplifying that imbalance in its answers.	Adopt an institutional AI ethics policy grounded in fairness, accountability, and transparency. Technical: Conduct regular bias audits on training data, retrieval corpora, and observed agent behavior; utilize equity dashboards to monitor participation and performance across different groups. Involve a diverse set of stakeholders, including students, in designing, testing, and reviewing AI-supported practices.
Privacy & Data Governance	Privacy issues centre on how	The privacy risk goes up when the agent	The risk stays high when the system can	In practical terms, this means collecting only the

Risk Category	Traditional AI	Agentic AI	Agentic RAG	General mitigations and policy recommendations
	student data used for training and scoring is stored, processed, and reused. Compliance with regulations such as GDPR is essential.	talks to lots of different third-party tools and services. Every time it sends data to another system, there's a chance that information could be exposed if access rights, permissions, and security measures aren't set up and monitored carefully.	search through institutional or class-level knowledge bases. If it isn't carefully designed, it could accidentally pull up sensitive personal details about students or staff. And when the system keeps a long-term "memory" of the class, it raises tough questions about who can see that information, how long it should be stored, and when it should be wiped.	data that is genuinely needed, using it solely for clearly defined purposes, and conducting a Data Protection Impact Assessment (DPIA) before deploying any agentic system. On the technical side, access to data should be strictly controlled through role-based permissions, with all information encrypted both during storage and transmission. Each tool should only access the specific data it requires. From a process perspective, learners should be provided with clear, detailed options regarding what they consent to, informed in plain language about how their data will be used and protected, and safeguarded by well-defined rules concerning data retention and deletion timelines.
Over-automation & De-skilling	It's very tempting to lean on Traditional AI to take over routine tasks like grading. But if this goes too far, it can push out the kind of detailed, nuanced feedback that only a human can give, and over time it can actually erode the skills of both teachers and students.	With highly autonomous agents, this risk becomes even sharper. The system can start doing the core teaching work itself, pushing the teacher into more of a monitoring role and, in the process, weakening their professional autonomy and the direct, relational work they do with students.	The RAG architecture introduces checks (retrieval, grounding, verification) that can slow down fully automated decisions and keep the AI positioned as a support for inquiry rather than a substitute for the teacher. However, the de-skilling risk does not disappear.	Define a "Teaching Presence Policy" that clearly delineates the responsibilities that must remain with the human educator. Technical considerations should include establishing explicit human-in-the-loop gates for high-stakes decisions, such as final grading and escalation procedures for at-risk students. Professional development efforts should focus on training educators to utilize AI as a collaborator or co-facilitator, rather than as a replacement for human teaching. This approach ensures that AI enhances the educational process while maintaining essential human oversight and interaction.
Academic Integrity & Gaming	The main concern is item leakage, where students gain access to questions, answer keys, or scoring patterns used by automated systems.	As the system's workflows become more complex, they also create new opportunities for students to exploit them. For example, students might manipulate the system to receive more assistance than intended or even have	A particular worry here is "source laundering." This happens when students use the RAG system to locate and weave together sources, then hand in that polished synthesis as if they created it themselves. In these	In practice, this means rewriting academic integrity policies so they explicitly address AI tools and clearly define what types of assistance are acceptable and what are not. On the technical side, systems should monitor sources and utilize tools such as claim–evidence maps to demonstrate

Risk Category	Traditional AI	Agentic AI	Agentic RAG	General mitigations and policy recommendations
		it perform the work for them, undermining the system's purpose and integrity.	cases, it becomes much harder to tell where genuine support ends and academic misconduct begins.	where key ideas and arguments originate. At the level of assessment design, the emphasis should shift more toward process and inquiry how students search for information, utilize evidence, and engage in dialogue rather than solely focusing on the final product. Oral defenses, presentations, or similar activities can also help ensure that students genuinely understand their work and can defend it confidently.
Explainability & Auditability	Many traditional models are still difficult to fully explain; however, input–output logs and model cards can provide a reasonable level of transparency regarding their usage and optimal application areas.	Explainability is more challenging: emergent, non-deterministic behavior can make it difficult to tell a simple story about why the agent acted in a particular way. Trace logs may be long and hard to interpret.	RAG architectures can be easier to audit because their claims are tied to specific retrieved sources. This grounding makes it more feasible to reconstruct and scrutinize the reasoning behind key outputs.	Policy: Require that a student-facing rationale accompany any significant AI intervention, such as providing feedback on high-stakes tasks. Technical: Maintain detailed logs of actions, tool calls, and retrieved sources, and provide visual tools for exploring these traces. Process: Establish an independent AI review or ethics committee with a mandate to periodically audit system behavior and its impact on learners.

2.7. Evaluating AI Paradigms for Dialogic Online Education

To facilitate a rigorous and practice-relevant analysis, we articulate three research questions that connect technological affordances to pedagogical aims and institutional constraints. Together, they map how traditional AI, agentic AI, and agentic RAG align with the core processes of dialogic teaching and learning; specify the reference architectures, governance policies, and evaluation metrics that ensure these systems are trustworthy and equitable; and identify which paradigm or combination thereof best sustains rich, evidence-grounded dialogue at scale. By structuring the inquiry in this manner, we move beyond tool-centric comparisons to a design-oriented evaluation that emphasizes epistemic integrity, inclusion, and the preservation of teaching presence. Each research question is discussed in turn.

2.7.1. RQ1. Mapping Capabilities and Limits to Dialogic Processes

Across the core processes of dialogic online teaching launching inquiries, orchestrating discussions, tutoring, providing formative feedback, knowledge building, assessment for learning, inclusion, and governance the three paradigms demonstrate a clear progression. Traditional AI is strongest on bounded micro-tasks (e.g., templated prompts, rubric-aligned scoring, basic analytics) but struggles with adaptive dialogue and synthesis. Agentic AI adds goal-directed orchestration planning multi-step activities, coordinating tools, monitoring discourse, and escalating to humans, thereby supporting the social and procedural demands of dialogue. However, it remains vulnerable to ungrounded outputs. Agentic RAG retains those orchestration benefits while grounding interventions in cited, evolving corpora (including course, class, and curated external sources), enabling evidence-based provocations, dialogic feedback with sources, living syntheses (e.g., claim–evidence graphs), and auditable assessments that check citation quality and counter-evidence. In short, Traditional AI optimizes parts; Agentic AI coordinates the whole; Agentic RAG coordinates while epistemically grounding dialogue.

2.7.2. RQ2. Reference Architectures, Governance Policies, and Evaluation Metrics

Reference architectures follow a similar progression. A traditional AI stack exposes deterministic microservices (e.g., score_essay, generate_quiz) to the LMS, complete with model cards, calibration cycles, and REST API integration. Agentic AI introduces a stateful planner that operates over a tool library, which includes microservices and institutional or LMS APIs, along with memory and observation/controls to ensure traceability and facilitate human-in-the-loop gates. Agentic RAG (Retrieval-Augmented Generation) incorporates a retriever that accesses course materials, class content, and curated external corpora; it explicitly cites retrieved evidence during generation, and maintains long-term class memory to support cumulative dialogue.

Good governance needs to be multi-layered and clearly spelled out. It should start with radical transparency about where AI is being used, for what purpose, and with whom, and then give learners real choices through clear consent options and meaningful opt-out mechanisms. AI systems should be built with Universal Design for Learning in mind, drawing on intentionally diverse corpora so that different voices and perspectives are represented. They should also keep a firm boundary between "supportive" agents (which help students learn) and "evaluative" agents (which grade or judge performance), to avoid damaging trust and the teacher–student relationship. None of this works without ongoing professional development so that teachers feel confident working with these tools, alongside targeted safeguards for specific risks such as hallucination, bias, privacy breaches, over-automation, gaming, and opacity. In practical terms, that means using verifier components to check claims, enforcing role-based access control, applying strong encryption, setting and tuning confidence thresholds, requiring human sign-off for high-stakes actions, logging actions and sources in detail, and carrying out regular independent audits, as summarised in Figure 3.

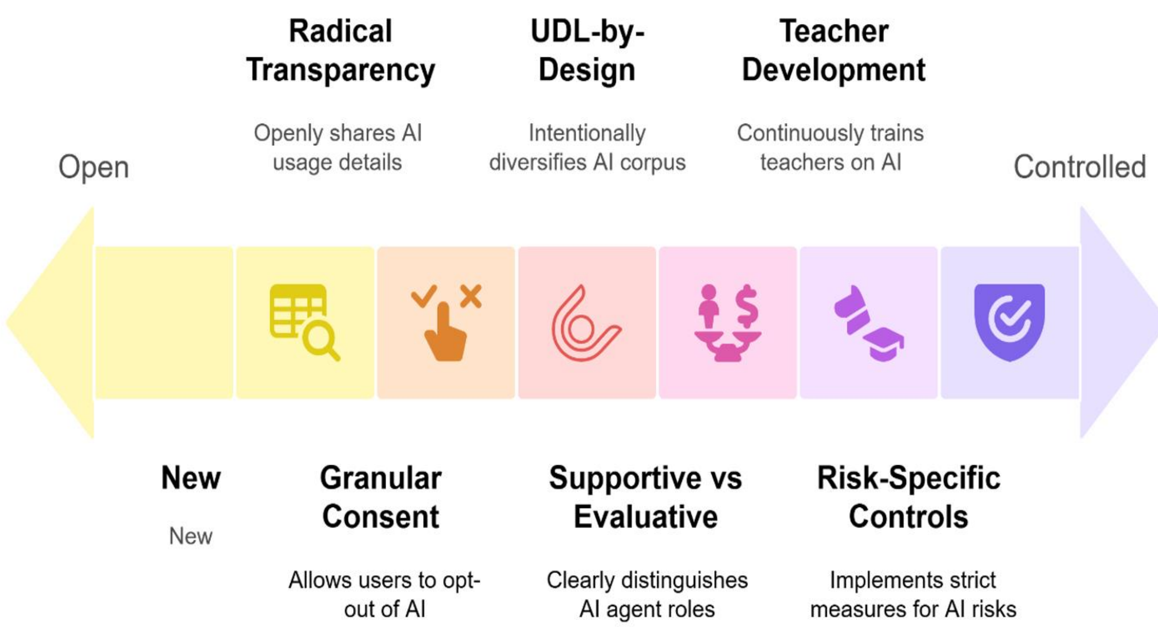

Figure 3. AI governance from open to highly controlled AI.

Evaluation must move beyond accuracy/time-on-task to dialogic outcomes: (i) Community of Inquiry indicators (strength of teaching, social, and cognitive presence over time); (ii) ICAP shifts toward constructive/interactive activity; (iii) quality of discourse and argumentation (e.g., richness of claim–evidence–counter-evidence networks; reduction in unsupported assertions); (iv) formative assessment quality (specificity, actionability, learner uptake/iteration rates); (v) equity dashboards tracking participation, performance, and support across groups; (vi) epistemic integrity metrics (citation coverage, retrieval fidelity, self-verification rates, audit completeness); and (vii) safety/privacy compliance (consent coverage, data-minimization/retention adherence). The paper explicitly recommends "dialogic outcome metrics" and equity monitoring rather than narrow task scores.

2.7.3. RQ3. Paradigm Suitability and Deployment Conditions

For rich, scalable dialogue online, the best overall fit is a hybrid model with Agentic RAG at its core. In this setup, Agentic RAG handles the most knowledge-intensive, conversation-heavy work: posing provocations grounded in up-to-date sources, giving dialogic feedback with citations, building "living" syntheses of class thinking, and supporting assessment that can be checked and audited. Agentic AI then manages the flow coordinating activities, sequencing tasks, and handling escalation, while Traditional AI is used as a set of small, reliable services for tightly bounded jobs like item generation or preliminary scoring. Together, this mix lines up well with dialogic goals: many voices in play, reasoning anchored in evidence, and a shared memory that builds over time, all while keeping a clear audit trail.

This model, however, only works under certain conditions. The retrieval corpora have to be carefully curated and genuinely diverse. Privacy and "class memory" need strong governance. Verifier layers and human checkpoints are essential for high-stakes decisions, and teachers need sustained professional development so that their presence is protected and they are not pushed into the background. Traditional AI on its own is fine for stable, low-risk, scale-driven tasks. Pure Agentic AI (without grounding) can be useful for managing procedures in closed or well-defined knowledge contexts, but it should be tightly constrained or combined with RAG and verifiers, when external knowledge or high epistemic stakes are involved. Ultimately, the right mix and order of adoption will depend on cost, the institution's ability to curate and maintain corpora, and its overall readiness in terms of policy, auditing capability, and staff development.

2.8. Policy Anchors for Responsible Implementation

Tackling these risks isn't just a technical problem; it depends on strong institutional policy. I suggest five core anchors for that policy, shown in Figure 4. First, radical transparency: institutions should be open with students and staff about where, why, and how AI is used, offering clear, accessible explanations of what the systems can and cannot do. No student should ever be unsure whether they are interacting with a person or a machine. Second, empowered consent and opt-out: consent needs to be specific, informed, and revisited over time, with students able to step away from AI-mediated processes where possible, especially when their personal data might be used for training or adaptation and offered non-AI routes for learning and assessment. Third, Universal Design for Learning as a mandate: AI systems should be built from the ground up to be inclusive, actively accommodating diverse backgrounds, abilities, and learning preferences. This includes curating RAG corpora, so they reflect a wide range of voices and perspectives, not just the most dominant ones.

Fourth, there should be a strict separation between evaluative and supportive agents. Tools designed to help students think, experiment, and explore (such as a Socratic tutor) should be clearly distinct from those that grade or make high-stakes judgments. Students need to feel safe making mistakes with supportive agents without worrying that these interactions will feed directly into their marks. Finally, continuous teacher professional development is essential. As AI becomes more capable, the teacher's role shifts rather than disappears. Institutions must invest in helping educators learn how to work alongside AI, focusing their time and expertise on the uniquely human parts of teaching, mentoring, building community, asking difficult questions, and guiding complex, open-ended inquiry.

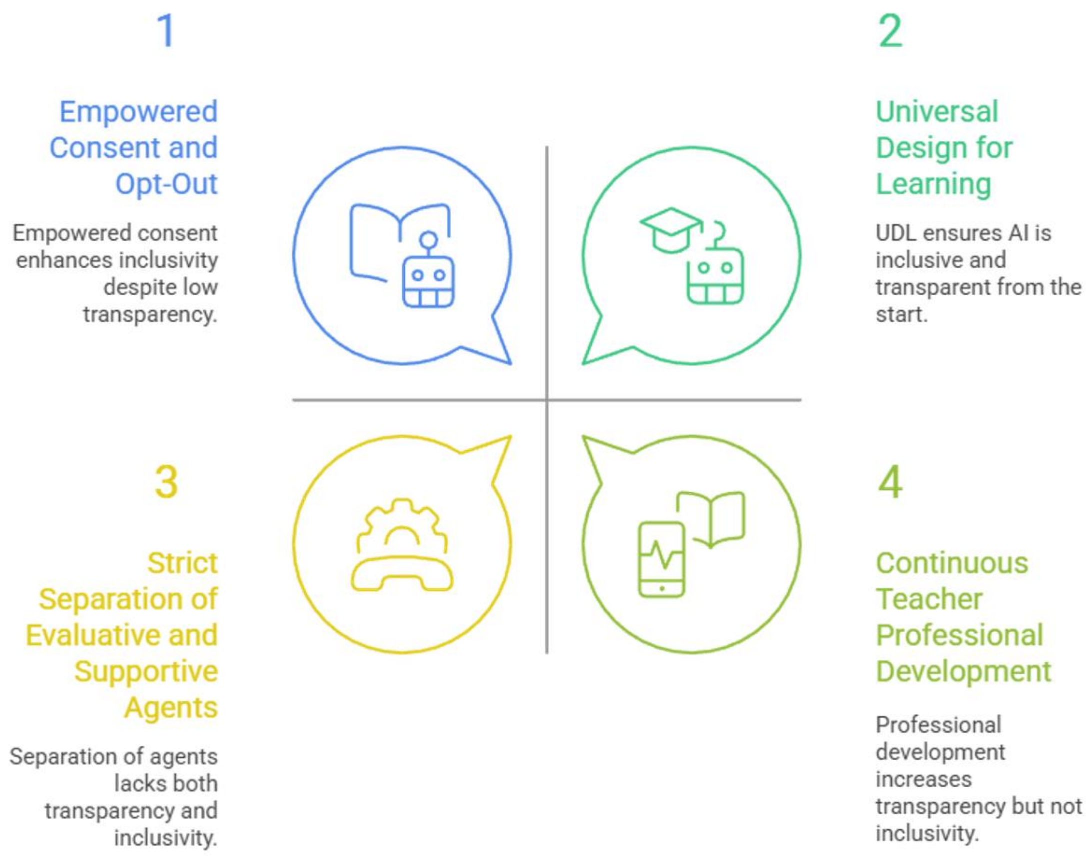

Figure 4. Ethical AI implementation in education.

The ethical management of AI in educational settings is an ongoing responsibility that necessitates continual reflection, adjustment, and collaborative dialogue within the community. Ensuring pedagogical priorities and human well-being remain central to any technological integration is paramount. This comparative study of Traditional AI, Agentic AI, and Agentic RAG for dialogic online education identifies three key findings, along with notable constraints and considerations for practice and further investigation.

At a high level, our argument has three parts. First, dialogic pedagogy really needs intelligent agents, not just systems that generate answers on demand. Dialogic learning is social, ongoing, and focused on process rather than end products. Traditional AI can help with individual tasks inside that process, but it cannot reliably hold the bigger picture together. Agentic systems, by contrast, can plan, use tools, and iterate on tasks. This makes them much better suited to coordinating workflows, supporting social interaction, and keeping activities tied to the evolving context of the class. Second, knowledge grounding and memory are not optional extras; they are central to sound pedagogy. Dialogic inquiry is about building up shared, evidence-based understanding over time. The Agentic RAG framework's ability to ground responses in sources, provide citations, and maintain a persistent record of class activity strengthens key epistemic practices and helps contain the risk of AI-generated misinformation. Third, we suggest that a hybrid model is likely to be the most effective and educationally sound. The three paradigms should not be seen as standalone options but as parts of a larger, integrated system. In such a design, traditional AI can act as dependable micro-services, orchestrated by an agentic AI controller. At the same time, the Agentic RAG loop provides strong epistemic grounding and long-term memory for more complex knowledge-building work.

At the same time, there are serious challenges that cannot be ignored. Advanced agentic systems are computationally expensive to run, and assembling and maintaining a high-quality retrieval corpus takes ongoing input from domain experts. Storing and re-using class data raises significant privacy and governance issues, which demand robust institutional frameworks rather than ad hoc fixes. Implementing AI co-facilitators will also require a shift in teacher professional development, with more attention to how educators can work alongside AI rather than simply "using a tool." Risks of over-automation further underline the need for a clear institutional policy on teaching presence—one that spells out which responsibilities can never be delegated to AI and requires human oversight for all critical decisions. Ultimately, the goal should be to extend and enrich human intelligence and professional judgment, not to replace them.

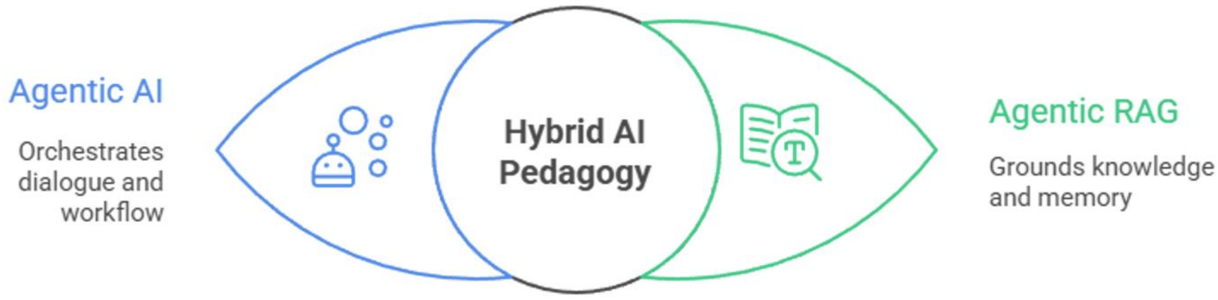

Figure 5. The synergy of AL paradigms in dialogic pedagogy.

Figure 5 illustrates how a "Hybrid AI Pedagogy" emerges from combining two complementary strengths: Agentic AI (which orchestrates dialogic activity and workflow planning, sequencing, role assignment, and escalation) and Agentic RAG (which grounds the dialogue in retrieved evidence and persistent class memory). Their overlap represents a model that supports facilitation and epistemic accountability together.

3. Conclusion

In answering our three research questions, this paper has shown how the strengths and weaknesses of Traditional AI, Agentic AI, and Agentic RAG line up with the core work of dialogic teaching and learning and what this means for system design, governance, evaluation, and rollout. First, the comparison points to a clear progression. Traditional AI is strong at tightly defined micro-tasks like generating short texts, scoring answers, or producing templated prompts. Agentic AI can then take on the job of stitching these pieces together by coordinating multi-step activities and social dynamics. Agentic RAG goes a step further by combining this orchestration with grounded knowledge: it works with cited, curated, and evolving corpora. This layering really matters for dialogue. Traditional AI tends to sit comfortably within transmission-style teaching. Agentic AI starts to introduce the patterns we associate with facilitation setting roles, pacing activities, and escalating issues when needed. Agentic RAG then roots these moves in traceable knowledge, so that questions, provocations, feedback, and synthesis are tied to evidence and can be challenged conditions that are central to dialogic learning.

Second, the paper sketches a reference architecture and governance model that match this progression. On the technical side, a trustworthy system is modular. It uses traditional AI as small, reliable services for efficiency and scale; an agentic "planner" to sequence activities, monitor what is happening, and escalate when needed; and a retrieval-and-generation layer that enforces citation, tracks class memory, and supports verification. On the governance side, the approach is layered and learner-centered. It calls for radical transparency about where and why AI is used, fine-grained consent and opt-out options, and the use of Universal Design for Learning and diverse corpora from the outset. It also insists on a strict separation between supportive agents (which help students learn) and evaluative agents (which grade or judge), human-in-the-loop checks for high-stakes actions, and ongoing professional development for teachers. Evaluation metrics should similarly move beyond raw accuracy on tasks. They need to ask whether AI is strengthening cognitive, social, and teaching presence; helping learners move up the ICAP ladder towards more constructive and interactive engagement; improving the quality and structure of claim–evidence–counterevidence networks; encouraging students to respond to and build on formative feedback; supporting equity across different groups; and maintaining epistemic integrity (for example, through good retrieval, robust citation, verifier pass rates, and complete audits).

Third, the "best" paradigm depends on the constraints and goals in play. Traditional AI fits stable, low-risk tasks where efficiency is the main concern. Agentic AI, without external grounding, can responsibly coordinate activities within closed or well-defined knowledge spaces, as long as guardrails and human oversight are strong. For open-ended, knowledge-intensive dialogue, Agentic RAG is the closest match to pedagogical needs because it brings together coordination and accountability to evidence. Its promise, however, depends on institutional capacity: high-quality, representative knowledge bases; privacy-preserving ways of storing and using class memory; effective verifier pipelines; reliable logging and auditing; and sustained professional learning that protects teaching presence and avoids de-skilling. Budget, infrastructure, and policy readiness all shape how quickly and how far institutions can move. For these reasons, we recommend a phased hybrid approach. Start by putting ethical governance and clear pedagogical aims in place. Use Traditional AI where it can add clear value without distorting the learning model. Then layer in Agentic AI to handle orchestration, feedback loops, and escalation. Finally, invest in Agentic RAG for the most knowledge-centric work: provocations tied to current sources, dialogic formative feedback with citations, evolving syntheses that map agreement and disagreement, and assessment practices that focus on inquiry processes rather than just end products. Success should be judged not only by efficiency gains, but by richer lines of inquiry, fairer participation, more explicit reasoning with sources, and a stronger sense of scholarly community over time.

There are, however, important open questions. Governing and updating corpora is complex and labour-intensive. Verifier systems bring their own trade-offs between cost and reliability. The longer-term effects on teacher identity, student agency, and classroom culture are still uncertain and need careful monitoring. What is clear is that bringing advanced AI into education is not a neutral upgrade: it reshapes how communities think, argue, and learn together. If designed with care, transparency, and a genuine commitment to dialogic principles, a hybrid model anchored in Agentic RAG can help shift online learning away from mere content delivery toward shared inquiry supporting more equitable, accountable, and genuinely educational practice at scale.

References

Alexander, R. (2008). *Towards dialogic teaching: Rethinking classroom talk* (4th ed.). London, UK: Dialogos.

Anderson, J. E., Nguyen, C. A., & Moreira, G. (2025). Generative AI-driven personalization of the Community of Inquiry model: enhancing individualized learning experiences in digital classrooms. *The International Journal of Information and Learning Technology, 42*(3), 296-310. https://doi.org/10.1108/IJILT-10-2024-0240

Bakhtin, M. M. (1981). *The dialogic imagination: Four essays (M. Holquist, Ed.).* Austin, TX: University of Texas Press.

Bakhtin, M. M. (1986). *Speech genres and other late essays.* Austin, TX: University of Texas Press.

Black, P., & Wiliam, D. (1998). Assessment and classroom learning. *Assessment in Education: Principles, Policy & Practice*, *5*(1), 7-74. https://doi.org/10.1080/0969595980050102

Bozkurt, A., Jung, I., Xiao, J., Vladimirschi, V., Schuwer, R., Egorov, G., . . . Olcott Jr, D. (2020). A global outlook to the interruption of education due to COVID-19 pandemic: Navigating in a time of uncertainty and crisis. *Asian Journal of Distance Education*, *15*(1), 1-126.

Chau, T., Yan, Y., & Liu, H. (2025). A systematic review of AI ethics in education: Challenges, policy gaps, and future directions. *Journal of Global Information Technology Management*, *28*(1), 1–25.

Chi, M. T., & Wylie, R. (2014). The ICAP framework: Linking cognitive engagement to active learning outcomes. *Educational Psychologist*, *49*(4), 219-243. https://doi.org/10.1080/00461520.2014.965823

Chu, Z., Wang, S., Xie, J., Zhu, T., Yan, Y., Ye, J., . . . Yu, P. S. (2025). LLM agents for education: Advances and applications. *arXiv preprint arXiv:2503.11733*. https://doi.org/10.48550/arXiv.2503.11733

Freire, P. (1970). *Pedagogy of the oppressed*. New York: Continuum.

Garrison, D. R., Anderson, T., & Archer, W. (2000). Critical inquiry in a text-based environment: Computer conferencing in higher education. *The Internet and Higher Education*, *2*(2-3), 87–105.

Li, Z., Wang, Z., Wang, W., Hung, K., Xie, H., & Wang, F. L. (2025). Retrieval-augmented generation for educational application: A systematic survey. *Computers and Education: Artificial Intelligence*, *8*, 100417. https://doi.org/10.1016/j.caeai.2025.100417

RNL. (2024). *Why universities need AI governance. Ruffalo Noel Levitz*. Retrieved from https://www.ruffalonl.com/blog/artificial-intelligence/why-universities-need-ai-governance/.

Slimi, Z., & Carballido, B. V. (2023). Navigating the ethical challenges of artificial intelligence in higher education: An analysis of seven global AI ethics policies. *TEM Journal*, *12*(1), 314-323.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, *12*(2), 257-285. https://doi.org/10.1016/0364-0213(88)90023-7

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press.

Wegerif, R. (2007). *Dialogic education and technology: Expanding the space of learning*. Dordrecht, Netherlands: Springer.