Scholastic Seed Inc.
e-Publishing Aggregator & Periodical Mentor

www.gjeis.com
global Journal of ENTERPRISE INFORMATION SYSTEM

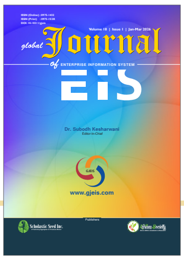

Navigating the Shift to the "Agentic AI University": Institutional Strategies for Operational and Pedagogical Resilience

– **Sudhansh Sharma***
Associate Professor, School of Computers and Information Sciences, IGNOU
sudhansh@ignou.ac.in *https://orcid.org/0000-0002-6870-0885*

ARTICLE HISTORY

Paper Nomenclature: Case Based Study (CBS)
Paper Code: GJEISV18I1JM2026CBS1
Submission at Portal (www.gjeis.com): 05-Jan-2026
Manuscript Acknowledged: 14-Jan-2026
Originality Check: 21-Jan-2026
Originality Test (Plag) Ratio (DrillBit): 03%
Author Reverts with Rectified Copy: 07-Feb-2026
Peer Reviewers Comment (Open): 15-Feb-2026
Single Blind Reviewers Explanation: 22-Feb-2026
Double-Blind Reviewers' Interpretation: 26-Feb-2026
Triple-Blind Reviewers' Annotations: 15-March-2026
Author Update (w.r.t. correction, suggestion & observation): 20- March-2026
Camera-Ready-Copy: 23-March-2026
Editorial Board Excerpt & Citation: 27-March-2026
Published Online First: 31-March-2026

ABSTRACT

Purpose: This paper examines how higher education institutions can respond to Agentic AI without reducing education to automation. It explores how autonomous or semi-autonomous AI agents are reshaping teaching support, assessment workflows, student services, analytics, and governance processes, while examining the implications for academic integrity, institutional accountability, data protection, and the cultivation of independent judgment.

Design/Methodology/Approach: The paper draws on recent discussions of autonomous agents in learning environments, the reported Einstein incident, policy guidance on ethical AI, and emerging evidence on cognitive effort. Based on these insights, it proposes a dual resilience model comprising operational resilience and pedagogical resilience to guide universities in responding to the opportunities and challenges of Agentic AI.

Findings: The study argues that operational resilience requires universities to treat AI as a governance concern rather than a narrow IT matter. This includes audit trails, human oversight, privacy safeguards, role-specific policies, and clearly protected manual-mode spaces. It further argues that pedagogical resilience requires a redesign of learning and assessment so that students learn to question AI outputs, defend their reasoning, connect work to lived experience, and remain responsible for the knowledge they present. The paper concludes that universities should not frame the issue as a choice between rejecting AI and adopting it uncritically, but should instead build systems in which AI improves access and support while human agency, ethical judgment, and academic credibility remain visible and enforceable.

Originality/Value: The paper proposes a dual resilience model that integrates operational resilience and pedagogical resilience as a framework for responding to the emergence of Agentic AI in higher education. It emphasizes balancing the benefits of AI-enabled support with the preservation of human agency, ethical judgment, and academic credibility.

Paper Type: Case Based Study

KEYWORDS: Agentic AI | Higher Education | AI Governance | Critical AI Literacy | Pedagogical Resilience | Assessment Integrity | Human Agency.

**Corresponding Author* (Sudhansh)

- Present Volume & Issue (Cycle): Volume 18 | Issue-1 | Jan-Mar 2026
- International Standard Serial Number:
  Online ISSN: 0975-1432 | Print ISSN: 0975-153X
- DOI (Crossref, USA) https://doi.org/10.18311/gjeis/2026
- Bibliographic database: OCLC Number (WorldCat): 988732114
- Impact Factor: 3.57 (2019-2020) & 1.0 (2020-2021) [CiteFactor]
- Editor-in-Chief: Dr. Subodh Kesharwani
- Frequency: Quarterly
- Published Since: 2009
- Research database: EBSCO https://www.ebsco.com
- Review Pedagogy: Single Blind Review/ Double Blind Review/ Triple Blind Review/ Open Review

- Publishers: Scholastic Seed Inc. and KARAM Society
- Place: New Delhi, India.
- Repository (figshare): 704442/13

Introduction

Universities have already spent several years adjusting to generative AI tools that can draft essays, summarize readings, translate text, prepare code, and answer questions in conversational form (Anderson, Watson, Rainie, & Anderson, 2025; Holmes et al., 2019). Those tools were disruptive, but they still worked mostly within a prompt-response pattern. A student, teacher, or administrator had to ask for something, read the output, and decide what to do next. Agentic AI changes that arrangement. An agent can be given a goal, break the goal into steps, interact with external systems, monitor outcomes, and revise its route when something fails (Kharbach, 2026a). In a university setting, such an agent may move through a Learning Management System, interpret instructions, prepare submissions, schedule tasks, retrieve resources, and communicate across platforms with limited human intervention.

This shift matters because the university is built around assumptions of participation and responsibility. When a student submits an assignment, attends an online session, completes a quiz, or asks for academic support, the institution usually assumes that the student is meaningfully involved. Agentic systems complicate that assumption. The reported Einstein incident, in which an autonomous tool was described as logging into Canvas, following lectures, completing course tasks, and submitting work, made the concern easy to understand (Kharbach, 2026a). The issue was not simply that a text had been generated by AI. The deeper concern was that the visible signs of student participation could be simulated by an agent. If the system can appear to attend, study, draft, and submit, then universities must reconsider what counts as evidence of learning.

The term "Agentic AI University" is used here to describe a higher education institution in which AI agents are not peripheral tools but embedded participants in academic and operational routines. Such a university may use agents to support advising, answer student queries, organize administrative work, monitor risk, personalize resources, and assist teaching staff (Kharbach, 2026a; Kharbach, 2026g). These uses can be valuable, especially for institutions facing large enrolments, uneven student support, and limited faculty time. However, the same infrastructure can also intensify surveillance, weaken student autonomy, over-standardize learning, and blur the boundary between assistance and substitution. The institutional challenge is therefore not whether AI should exist on campus; it already does. The challenge is how to govern it without losing the human purposes that justify higher education in the first place (Anderson & Rainie, 2026; UNESCO, 2023).

This paper develops that challenge through two linked forms of resilience. Operational resilience refers to the capacity of universities to manage agentic AI responsibly across policies, data systems, accountability structures, and day-to-day workflows (U.S. Department of Education, Office of Educational Technology, 2024). Pedagogical resilience refers to the capacity of teaching and assessment systems to preserve real learning when students have access to tools that can produce polished work, solve problems, and navigate academic tasks (Kharbach, 2026d; Lee et al., 2025). These two forms of resilience cannot be separated. A well-designed assessment model will fail if institutional systems cannot verify responsibility, protect data, or explain decisions. Similarly, strong policies will have limited value if classroom practice continues to reward only the final product rather than the thinking that produced it.

Table 1. Key Differences between Generative AI and Agentic AI in Higher Education

Dimension	Generative AI tools	Agentic AI systems
Primary action	Respond to a user prompt and produce text, images, code, or summaries.	Pursue a goal through multiple steps, tools, and decisions.
Typical university use	Drafting, tutoring, brainstorming, feedback, translation, and content support.	Navigating LMS platforms, organizing tasks, coordinating support, submitting forms, and managing workflows.
Human involvement	The user normally initiates each exchange and checks the output.	The user may set the objective while the system handles much of the execution.
Main academic risk	Unacknowledged authorship, shallow writing, and over-reliance on generated content.	Substitution of student participation, automated impersonation, and opaque decision chains.
Governance need	Guidance on acceptable use, citation, and assessment redesign.	Auditability, access controls, human oversight, boundary-setting, and incident response.

Table 1 clarifies why the agentic turn requires a stronger institutional response than earlier debates about chatbots. A chatbot can certainly be misused, but an autonomous agent can connect several actions into a workflow (Kharbach, 2026a). This creates a different class of risk. It also requires a different class of policy, because the university must govern not only what a student writes but also what digital actions are taken on the student's behalf, by whom, under what permissions, and with what record of accountability (U.S. Department of Education, Office of Educational Technology, 2024).

The paper is organized as follows. The literature review discusses the movement from generative to agentic AI, the governance implications for universities, the cognitive risks associated with passive AI use, and the need for critical AI literacy. The research gaps section identifies areas where current scholarship remains thin. The methodology section proposes a mixed-method design suitable for studying institutional and pedagogical responses to agentic AI. The results and discussion section synthesizes likely findings from institutional audits, discipline-based pilots, assessment redesign, and participatory policy work. The conclusion presents a practical direction for universities seeking to remain credible, inclusive, and human-centred in an AI-rich environment.

Literature Review

The literature on AI in higher education has developed quickly, but much of the early discussion focused on generative systems that respond to prompts. These systems affected academic writing, feedback, tutoring, and lesson preparation (Holmes et al., 2019; Kharbach, 2025). Agentic AI extends the problem because it is organized around action rather than response alone. It can decide what needs to be done next, call tools, use interfaces, and continue a task after the first output (Kharbach, 2026a). In education, this creates both opportunity and unease. A well-designed agent could help a first-generation student locate support services, remind a learner about deadlines, or guide a teacher through repetitive administrative tasks. A poorly governed agent could also complete academic routines in ways that hide the absence of genuine student engagement.

The Einstein example has become a useful reference point because it illustrates the difference between assisted production and automated participation (Kharbach, 2026a). Universities have experience dealing with ghostwriting, plagiarism, and contract cheating, but agentic AI makes the problem faster, cheaper, and harder to observe. If a system can log in, follow instructions, generate responses, and submit work, then the institution must ask whether attendance logs, submission timestamps, and written outputs still provide reliable evidence of learning. This does not mean that all digital learning must be distrusted. It means that learning systems require stronger evidentiary design, including process records, human confirmation, oral explanation, and assessment tasks that are linked to context and experience (UNESCO, 2023).

Another strand of the literature emphasizes governance. AI in universities is often first handled by IT units because it appears to be a tool, a platform, or a procurement issue. That framing is too narrow. Agentic systems can affect privacy, academic integrity, accessibility, labour expectations, student support, institutional liability, and the distribution of opportunity (Anderson & Rainie, 2026; U.S. Department of Education, Office of Educational Technology, 2024). For that reason, safe AI adoption requires the involvement of academic leaders, faculty boards, student representatives, ethics committees, data protection officers, accessibility teams, and legal advisors. Governance should not be reduced to a single acceptable-use statement. It should include decision rights, documentation requirements, escalation paths, review cycles, and procedures for addressing misuse or harmful outcomes.

Operational resilience is therefore more than technical robustness. It is the institution's ability to continue serving students fairly when AI tools fail, behave unpredictably, or are used in ways that were not intended (Anderson & Rainie, 2026). A resilient university maintains human oversight over consequential decisions, keeps audit trails for automated workflows, limits data collection to legitimate educational purposes, and explains how AI-supported decisions can be challenged (U.S. Department of Education, Office of Educational Technology, 2024; UNESCO, 2023). It also avoids the temptation to automate every friction point. Some friction is educationally useful. The pause before answering, the struggle to form an argument, the effort to read carefully, and the need to justify a claim are not inefficiencies to be removed; they are part of learning.

The literature on cognition adds a second concern. Researchers have described risks such as cognitive offloading, thinking atrophy, and metacognitive laziness (Fan et al., 2025; Lee et al., 2025). These terms point to a familiar pattern: when a tool supplies answers quickly, learners may spend less time planning, monitoring, checking, and revising their own thinking. This does not make AI inherently damaging. Calculators, search engines, and writing tools have also changed how students work. The difference is that generative and agentic AI can supply not only information but also structure, interpretation, language, and strategy. If students accept these outputs without examination, they may complete tasks while doing less of the mental work that builds durable understanding.

At the same time, the literature does not support a simple anti-AI position. Students can use AI in ways that increase learning when tasks require them to compare outputs, detect hallucinations, identify missing assumptions, rewrite explanations in their own words, and defend final choices (Kharbach, 2026d; The Open University, 2025). In this sense, the central issue is not AI use but the quality of the learning relationship created around AI. Passive use turns AI into a substitute for thought. Critical use turns AI into an object of inquiry. This distinction is essential for curriculum design. A university that only bans AI may push usage underground. A university that permits all AI use without structure may normalize intellectual dependency. The more productive route is to teach students when to use AI, when not to use it, and how to remain accountable for what they submit.

Critical AI literacy is one answer to this problem. It involves more than teaching students how to write prompts (The Open University, 2025). Students need to understand that AI outputs can be fluent without being true, confident without being justified, and convenient without being fair. They should learn to inspect sources, test claims, recognize bias, protect personal data, and disclose assistance honestly (UNESCO, 2023). Existential literacy extends the discussion further by asking what it means to act, choose, and learn in a world where machines can generate advice and apparent expertise (Anderson & Rainie, 2026). These questions may sound philosophical, but they have practical consequences. A student who has never had to judge the limits of AI may struggle to take responsibility in professional settings where automation affects real people.

Table 2. Student Motivational Needs in AI-Assisted Learning

Self-Determination Theory construct	Potential support from AI	Possible harm when AI is poorly framed
Autonomy	Students can access flexible guidance, choose learning routes, and review material at their own pace.	Students may feel that the system is directing their choices or doing the work for them.
Competence	Immediate feedback can help learners identify mistakes and improve drafts before class discussion.	Quick answers may hide weak understanding, especially when students cannot verify the output.
Relatedness	AI may reduce isolation for learners who hesitate to ask basic questions publicly.	Overuse may reduce peer interaction, teacher contact, and the social confidence built through dialogue.

Note: The table adapts the Self-Determination Theory lens to the agentic AI context and summarizes the motivational tension commonly reported in AI-assisted learning environments (Du & Alm, 2024).

Assessment has become the most visible pressure point. Conventional essays and take-home assignments still have educational value, but they can no longer carry the whole burden of proof. A polished written submission may show that a student can manage a tool, but it may not show what the student knows, remembers, questions, or can apply under dialogue (Kharbach, 2026c; UNESCO, 2023). This does not imply that universities should abandon writing. Instead, writing should be embedded in richer assessment sequences. For example, students may submit a draft history, explain their revisions, answer oral questions, connect the work to a practical case, or complete a short in-class transfer task. Such designs make authorship and understanding easier to observe.

The literature also suggests that assessment redesign must be realistic. Oral examinations, live demonstrations, and video reflections can strengthen academic integrity, but they require time, training, rubrics, and moderation (Kharbach, 2026c; Kharbach, 2026d). In large classes, poorly planned alternatives may create unfair workloads for teachers and anxiety for students. The goal should not be to make every task AI-proof. No assessment is permanently AI-proof. The better aim is to make assessment more learning-rich, transparent, and varied so that students cannot succeed by outsourcing the entire process to a machine.

Research Gaps

Several gaps remain in current research on agentic AI in higher education. The first gap concerns time. Many studies examine immediate reactions to AI tools or short-term classroom experiments. Much less is known about what happens when students rely on AI agents across a full degree programme. Longitudinal research is needed to understand effects on memory, writing confidence, problem-solving, professional identity, collaboration, and the willingness to persist through difficult learning tasks (Du & Alm, 2024; Smith et al., 2025).

The second gap concerns institutional diversity. A large share of existing discussion draws on well-resourced universities, technology-rich settings, or general student populations. Agentic AI may look very different in community colleges, open universities, rural institutions, vocational programmes, clinical training, teacher education, and universities in the Global South. These settings may face weaker infrastructure, different regulatory requirements, different language realities, and different student support needs. A governance model that works in one context may not transfer easily to another (UNESCO, 2023; U.S. Department of Education, Office of Educational Technology, 2024).

The third gap is practical governance. Many papers recommend responsible AI use, but fewer explain how a university can coordinate responsibility across departments. Institutions need models that connect procurement, data governance, teaching policy, academic integrity, disability support, staff development, student voice, cybersecurity, and appeals procedures (Anderson & Rainie, 2026; U.S. Department of Education, Office of Educational Technology, 2024). Without such integration, universities may produce attractive policy documents while day-to-day practices remain fragmented.

The fourth gap concerns equity. Agentic AI may widen inequalities if affluent students and institutions gain access to better tools, safer platforms, and stronger training while others rely on free, opaque, or less secure systems. Equity also includes linguistic and accessibility issues. Students with disabilities, multilingual learners, and students with limited digital access may experience both benefits and risks that differ from those of more privileged groups. Research must therefore examine who gains, who is exposed to harm, and who is asked to carry the cost of adaptation (UNESCO, 2023; U.S. Department of Education, Office of Educational Technology, 2024).

The final gap concerns assessment scalability. Educators often recommend oral defences, live demonstrations, reflective tasks, and process-based portfolios (Kharbach, 2026c; Kharbach, 2026d). These methods are promising, but large enrolments can make them difficult to sustain. More evidence is needed on rubrics, sampling models, peer moderation, group viva formats, and technology-supported workflows that protect integrity without overwhelming faculty.

Research Methodology

This paper proposes a mixed-methods design for studying institutional and pedagogical resilience in the Agentic AI University. The design is appropriate because agentic AI is not only a classroom matter. It cuts across governance, policy, teaching, assessment, student support, data systems, and organizational culture (Anderson & Rainie, 2026; U.S. Department of Education, Office of Educational Technology, 2024). A single survey or one-off experiment would not capture that complexity. The proposed methodology therefore combines institutional case studies, discipline-based learning pilots, assessment redesign trials, system stress-testing, and participatory policy development.

The first component consists of institutional case studies. Three types of institutions would be selected: a public research university, a professional college with practice-based programmes, and an online or blended institution. The aim is not to rank them but to understand how context shapes AI readiness. Data would be collected through interviews with senior administrators, faculty, IT staff, academic integrity officers, student support teams, and student representatives. Policy documents would also be reviewed, including AI guidelines, privacy notices, assessment policies, procurement rules, and codes of conduct. The analysis would examine whether the institution treats AI as a governance issue or only as a technology issue (Anderson & Rainie, 2026; U.S. Department of Education, Office of Educational Technology, 2024).

The second component examines learning behaviour across disciplines. Students from STEM, humanities, social sciences, and professional programmes would complete AI-assisted tasks under structured conditions. Some tasks would allow ordinary AI support, while others would require students to challenge outputs, document prompts, explain rejected suggestions, and produce a final reflective note. Data would be gathered through think-aloud sessions, short surveys, learning journals, and teacher observations. This component would help identify when AI supports learning and when it encourages shallow acceptance (Fan et al., 2025; Lee et al., 2025; The Open University, 2025).

The third component focuses on assessment. Selected courses would pilot layered assessment designs. Instead of relying only on written submissions, students would complete a sequence such as proposal, annotated source check, draft, oral explanation, applied scenario, and final reflection. Comparable courses would retain conventional assessment formats. The study would compare evidence of understanding, student workload, faculty workload, perceived fairness, and indicators of authorship. The intention is not to prove that one method is perfect, but to identify combinations that are defensible and manageable (Kharbach, 2026c; Kharbach, 2026d; UNESCO, 2023).

The fourth component is institutional stress-testing. Universities routinely test cybersecurity, but they less often test whether academic systems can withstand agentic misuse. A controlled red-team exercise could simulate how an agent might navigate an LMS, submit forms, produce routine academic work, or exploit weak authentication. The exercise would be conducted ethically, with permission, limited scope, and no exposure of student data. The purpose would be to identify weaknesses in access controls, audit logs, verification points, and incident response procedures (Kharbach, 2026a; U.S. Department of Education, Office of Educational Technology, 2024).

The fifth component involves participatory policy design. Faculty, students, administrators, accessibility specialists, and data protection officers would take part in workshops where they review findings and draft practical rules. These workshops would distinguish between AI-supported, AI-disclosed, AI-limited, and AI-prohibited activities. They would also identify manual-mode spaces where students must demonstrate skill without AI. This participatory element is important because rules imposed from the top may be ignored if they do not reflect classroom realities (Anderson & Rainie, 2026; UNESCO, 2023).

Qualitative data would be analysed through thematic coding, with attention to governance, agency, trust, workload, equity, and assessment integrity. Quantitative data from surveys and course comparisons would be analysed descriptively and, where sample size permits, through comparative tests. Reliability would be strengthened through triangulation across interviews, documents, student work, observations, and policy workshops. Ethical safeguards would include informed consent, data minimization, anonymization, secure storage, and clear boundaries on any stress-testing activity (UNESCO, 2023; U.S. Department of Education, Office of Educational Technology, 2024).

Results and Discussion

The expected findings from the proposed design point toward a consistent pattern: universities are beginning to recognize agentic AI as a major issue, but many remain more prepared at the level of classroom advice than institutional coordination. Interviews with academic leaders and staff would likely show that policy language is improving, especially around disclosure and academic integrity (Anderson et al., 2025; Kharbach, 2026b). However, the operational details often remain vague. Staff may know that AI use should be responsible, but they may not know who approves new tools, how student data are protected, how automated decisions are reviewed, or what to do when an autonomous agent appears to have completed work on behalf of a student.

Document analysis would likely reveal a second pattern. AI policies often speak in broad ethical terms but do not always connect to procurement, cybersecurity, disability accommodation, appeals, or staff workload (U.S. Department of Education, Office of Educational Technology, 2024). This creates a gap between principle and practice. For example, a policy may require transparency, but the LMS may not record enough detail to show whether a workflow was completed by a student, an assistant, or an agent. A policy may require human oversight, but staff may not have time or training to review AI-supported decisions. Operational resilience therefore depends on the boring but essential work of implementation: logs, permissions, review points, role descriptions, and accountability chains.

Table 3. Methodological Components and Evidence Sources

Component	Evidence collected	Purpose
Institutional case studies	Interviews, policies, governance documents, and workflow maps.	Identify how universities frame, control, and review agentic AI.
Discipline-based pilots	Think-aloud data, student reflections, surveys, and teacher observations.	Examine how AI affects reasoning, confidence, and critical engagement across fields.
Assessment redesign trials	Rubrics, oral responses, process portfolios, workload logs, and student feedback.	Test whether layered assessment provides stronger evidence of learning.
AI stress-testing	Controlled red-team notes, audit logs, access-control findings, and incident reports.	Locate weak points in institutional systems before misuse occurs.
Participatory policy workshops	Draft rules, stakeholder feedback, scenario responses, and implementation notes.	Produce practical guidance that is understood by the people expected to use it.

Table 4. Operational Readiness Indicators for Agentic AI Governance

Readiness area	Low-readiness signal	Resilient practice
Policy ownership	AI guidance is owned only by IT or written as a generic acceptable-use note.	Academic, technical, legal, accessibility, and student voices share responsibility.
Auditability	Automated actions leave limited records or cannot be linked to a responsible user.	High-risk workflows include logs, timestamps, permissions, and reviewable decision trails.
Human oversight	Human review is mentioned but not scheduled, staffed, or documented.	Clear checkpoints identify when a person must confirm, override, or explain a decision.
Data protection	Tools are adopted before data flows and retention rules are understood.	Data minimization, vendor review, consent, and retention limits are built into adoption.
Manual-mode spaces	Every academic task is assumed to be open to AI assistance.	Core skills are protected through in-person, oral, practical, or process-based demonstrations.

The discipline-based pilots would likely produce a more nuanced picture than a simple claim that AI either helps or harms learning. Students using AI passively may report lower effort in comprehension, synthesis, and drafting (Fan et al., 2025; Lee et al., 2025). Their work may appear more fluent but less personally grounded. In think-aloud sessions, such students may read AI responses with limited questioning, accept the first plausible answer, and struggle to explain why a claim is valid. This pattern supports the concern that AI can reduce productive struggle when tasks are designed around output alone.

By contrast, students who are required to challenge AI outputs may show stronger engagement. When the task asks them to identify hallucinations, compare sources, record disagreements, and defend final choices, the AI becomes less of a shortcut and more of a thinking partner (Kharbach, 2026d; The Open University, 2025). The difference lies in task design. A passive prompt asks the student to receive. A resilient prompt asks the student to inspect, argue, revise, and own the result. This finding is central to pedagogical resilience because it shifts attention from banning tools to designing learning conditions that make thought visible.

The assessment trial would likely confirm that traditional essays are vulnerable when used in isolation. Many submissions may share similar structures, vocabulary, and argument patterns, especially when students rely heavily on the same tools (Kharbach, 2026c; UNESCO, 2023). This convergence does not always prove misconduct, but it weakens the teacher's ability to distinguish independent understanding from automated assistance. Layered assessment offers a stronger alternative. When a student must explain a draft orally, apply an idea to a local case, respond to a follow-up question, or show how feedback changed the work, the teacher gains better evidence of learning.

The main difficulty is workload. Faculty may value oral questioning and process portfolios but find them hard to manage in large classes. This is where institutional support becomes necessary. Universities could use short sampling vivas, group-based defences, rotating oral checks, structured reflection templates, and teaching assistant training to make richer assessment more feasible (Kharbach, 2026c; U.S. Department of Education, Office of Educational Technology, 2024). Departments may also need to decide which courses require strict manual-mode demonstration and which can allow more AI-supported work. Not every assignment needs the same level of control. A first-year writing task, a capstone project, a clinical competency, and an administrative training module have different risks.

The stress-testing component would likely expose the difference between platform security and educational assurance. A system may be technically secure in the sense that it requires a password, yet still educationally weak if an agent can operate after login without meaningful checks (Kharbach, 2026a). Institutions may therefore need additional verification for high-stakes actions, such as identity confirmation, oral defence triggers, unusual activity alerts, and clear records of delegated access. These safeguards must be proportionate. Excessive surveillance can harm trust and privacy. The aim should be targeted assurance, not permanent monitoring of every student action (UNESCO, 2023).

Equity would appear across all findings. Some students may benefit greatly from AI support, particularly those who need help organizing tasks, understanding dense readings, or practising language. Others may be disadvantaged if they cannot access reliable tools or if institutional rules are unclear (UNESCO, 2023; U.S. Department of Education, Office of Educational Technology, 2024). Faculty equity also matters. If assessment redesign increases workload without support, the burden will fall unevenly on already stretched teachers. A resilient university therefore considers access, training, infrastructure, and workload as part of the same problem.

The discussion leads to a practical principle: universities should move from product policing to process assurance. Product policing tries to decide whether a finished text was produced by a human or a machine. That task is becoming less reliable. Process assurance asks how the work was developed, where human judgment appeared, what support was used, and whether the student can explain, defend, and apply the knowledge (Kharbach, 2026c; Lee et al., 2025). This approach is not perfect, but it is more educationally meaningful. It also reduces over-dependence on detection tools, which can produce uncertain results and may treat polished or formulaic academic writing suspiciously.

A second principle is cognitive triage. Universities should decide which cognitive activities can be safely supported by AI and which should be protected for human practice (Fan et al., 2025; Lee et al., 2025). For example, AI may help students generate practice questions, translate instructions, or receive early feedback. However, students still need opportunities to read without summaries, solve problems without automated hints, speak without scripts, and make judgments without machine recommendations. Cognitive triage does not reject AI; it places boundaries around the skills that education is meant to build.

A third principle is shared accountability. Students should disclose meaningful AI assistance, but responsibility cannot rest only with students. Teachers must design tasks that make expectations clear. Departments must provide assessment policies that are realistic. Administrators must review tools before adoption. Vendors must be held to standards of privacy and explainability. Universities must offer training rather than assuming that staff and students will interpret AI rules correctly on their own (Anderson et al., 2025; U.S. Department of Education, Office of Educational Technology, 2024). Agentic AI turns responsibility into a networked matter, and governance must reflect that network.

Taken together, the findings suggest that the Agentic AI University should not be imagined as a fully automated campus. A better model is a deliberately bounded university: one that uses AI for support, access, and efficiency, but also preserves human-only moments where students must think, explain, struggle, and decide (Anderson & Rainie, 2026; The Open University, 2025). These moments are not nostalgic attachments to older education. They are the conditions under which trust in learning can survive.

Recommendations

First, universities should establish an AI governance board or equivalent cross-functional body with authority to review high-risk tools, data flows, assessment implications, accessibility issues, and student concerns. This body should not operate as a symbolic committee. It should maintain a living register of approved systems, risk levels, review dates, and responsible owners (U.S. Department of Education, Office of Educational Technology, 2024).

Second, institutions should classify academic tasks according to permitted AI involvement. A simple four-level model may be used: AI-prohibited tasks, AI-limited tasks, AI-disclosed tasks, and AI-integrated tasks. This classification should appear in course outlines and assessment rubrics so that students do not have to guess what is allowed (Kharbach, 2026c; UNESCO, 2023).

Third, critical AI literacy should be embedded across the curriculum. Students should practise source checking, hallucination detection, prompt critique, bias analysis, disclosure writing, and oral defence of AI-assisted work. These skills should not be confined to computer science or digital literacy courses because agentic AI affects all fields (Kharbach, 2026d; The Open University, 2025).

Fourth, universities should create manual-mode spaces for core competencies. These may include in-class writing, viva voce checks, studio demonstrations, lab-based problem solving, clinical simulations, field reflections, or live case analysis. The purpose is to preserve evidence of individual competence, not to punish students for using technology elsewhere (Anderson & Rainie, 2026; UNESCO, 2023).

Finally, assessment reform must be supported with workload planning. Faculty should receive templates, moderation guidance, sample rubrics, time allowances, and technical support. Without such support, even well-intentioned reform may fail in practice (Kharbach, 2026c; U.S. Department of Education, Office of Educational Technology, 2024).

Conclusion

Agentic AI marks a structural shift in higher education because it can act across systems rather than merely generate responses (Kharbach, 2026a). This development creates genuine opportunities for student support, administrative efficiency, and personalized learning. It also threatens academic integrity, student agency, privacy, and the reliability of traditional assessment signals (UNESCO, 2023; U.S. Department of Education, Office of Educational Technology, 2024). The central question for universities is therefore not whether AI should be used, but how its use can be governed and taught in ways that preserve the meaning of education.

This paper has argued for a dual resilience model. Operational resilience requires universities to build policies, audit systems, human oversight, data safeguards, and accountable workflows (Anderson & Rainie, 2026; U.S. Department of Education, Office of Educational Technology, 2024). Pedagogical resilience requires teachers to design learning and assessment that make reasoning visible, require students to challenge AI outputs, and protect spaces for unaided demonstration of core skills (Kharbach, 2026d; The Open University, 2025). These forms of resilience are mutually dependent. A university cannot teach responsible AI use without responsible systems, and it cannot govern AI credibly if classroom practice continues to ignore the realities of agentic tools.

The Agentic AI University should not be a campus where machines quietly replace human participation. It should be a university that uses intelligent systems carefully while strengthening the human capacities that education exists to develop: judgment, responsibility, creativity, dialogue, ethical awareness, and the ability to learn with others. If institutions build those values into governance and pedagogy now, agentic AI can become a managed transformation rather than an uncontrolled disruption (Anderson & Rainie, 2026; UNESCO, 2023).

References

- AI for Education, & Student Achievement Partners. (2025). Guide to integrating generative AI for deeper math learning. https://www.aiforeducation.io/ai-resources/guide-to-integrating-generative-ai-for-deeper-math-learning
- Anderson, D. J., Watson, C. E., Rainie, L., & Anderson, J. (2025). Student guide to artificial intelligence. Elon University; American Association of Colleges and Universities. https://studentguidetoai.org/wp-content/uploads/2025/03/Student-Guide-to-AI-2025.pdf
- Anderson, J., & Rainie, L. (2026, April). Building a human resilience infrastructure for the AI age. Imagining the Digital Future Center, Elon University. https://imaginingthedigitalfuture.org/reports-and-publications/human-resilience-in-the-age-of-ai/
- Awadallah Alkouk, W., & Khlaif, Z. N. (2024). AI-resistant assessments in higher education: practical insights from faculty training workshops. https://www.researchgate.net/publication/386384662_AI-resistant_assessments_in_higher_education_practical_insights_from_faculty_training_workshops
- Du, J., & Alm, A. (2024). The impact of ChatGPT on English for Academic Purposes (EAP) students' language learning experience: A Self-Determination Theory perspective. Education Sciences, 14(7), 726. https://doi.org/10.3390/educsci14070726. https://www.mdpi.com/2227-7102/14/7/726
- Fan, Y., Tang, L., Le, H., Shen, K., Tan, S., Zhao, Y., Shen, Y., Li, X., & Gašević, D. (2025). Beware of metacognitive laziness: Effects of generative artificial intelligence on learning motivation, processes, and performance. https://bera-journals.onlinelibrary.wiley.com/doi/epdf/10.1111/bjet.13544
- Holmes, W., Bialik, M., & Fadel, C. (2019). Artificial intelligence in education: Promises and implications for teaching and learning. Center for Curriculum Redesign. https://www.researchgate.net/publication/332180327_Artificial_Intelligence_in_Education_Promise_and_Implications_for_Teaching_and_Learning
- Kharbach, M. (2025). The AI-ready teacher: Strategies for building adaptive expertise. Educators Technology. https://medkharbach.com/wp-content/uploads/2026/01/AI-Tools-and-Resources-Your-Companion-for-Integrating-AI-in-the-classroom-for-colleagues.pdf
- Kharbach, M. (2026a). AI for Teachers in 2026, What worth using and Why? https://medkharbach.com/wp-content/uploads/2026/01/AI-for-Teachers-in-2026.pdf
- Kharbach, M. (2026b). Teaching With AI: A Step by Step Guide for Teachers. https://www.educatorstechnology.com/2025/02/teaching-with-ai.html
- Kharbach, M. (2026c). AI Tools for Teachers - Explore our curated collection of educational aitools for teachers and educators. https://medkharbach.com/wp-content/uploads/2026/01/AI-for-Teachers-in-2026-1.pdf
- Kharbach, M. (2026d). Critical thinking activities for the age of AI: Hands-on exercises for faculty and students. Educators Technology. https://www.educatorstechnology.com/wp-content/uploads/2026/04/critical-thinking-activities-for-the-age-of-AI.pdf
- Kharbach, M. (2026e). Microsoft Copilot for education: A teacher's complete guide. Educators Technology. https://www.educatorstechnology.com/wp-content/uploads/2026/04/Microsoft-Copilot_Education_Guide.pdf?trk=public_post_comment-text
- Kharbach, M. (2026f, April). The AI activities guide for teachers: Hands-on strategies and tools for your classroom. Educators Technology. https://www.educatorstechnology.com/wp-content/uploads/2026/04/AI-Activities-Guide-for-Teachers.pdf
- Kharbach, M. (2026g). The educator's guide to AI tools: Practical tools for lesson planning, feedback, assessment, and beyond. Educators Technology. https://www.educatorstechnology.com/wp-content/uploads/2026/04/Educators-Guide-to-AI.pdf?trk=public_post_comment-text
- Lee, H.-P., Sarkar, A., Tankelevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025). The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and mastery. Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '25). https://doi.org/10.1145/3706598.3713778. https://www.researchgate.net/publication/391270185_The_Impact_of_Generative_AI_on_Critical_Thinking_Self-Reported_Reductions_in_Cognitive_Effort_and_Confidence_Effects_From_a_Survey_of_Knowledge_Workers
- Smith, J. M., Dukes, J., Sheldon, J., Nnamani, M. N., Esteves, N., & Reich, J. (2025). A guide to AI in schools: Perspectives for the perplexed. MIT Teaching Systems Lab. https://tsl.mit.edu/ai-guidebook/
- https://tsl.mit.edu/wp-content/uploads/2025/08/GuideToAIInSchools.pdf
- The Open University. (2025). A framework for the learning and teaching of critical AI literacy skills (Version 0.1). https://www.open.ac.uk/blogs/learning-design/wp-content/uploads/2025/01/OU-Critical-AI-Literacy-framework-2025-external-sharing.pdf
- United Nations Educational, Scientific and Cultural Organization. (2023). Guidance for generative AI in education and research. https://doi.org/10.54675/EWZM9535. https://www.unesco.org/en/articles/guidance-generative-ai-education-and-research
- U.S. Department of Education, Office of Educational Technology. (2024, October). Empowering education leaders: A toolkit for safe, ethical, and equitable AI integration. https://files.eric.ed.gov/fulltext/ED661924.pdf

GJEIS Prevent Plagiarism in Publication

DELNET-Developing Library Network, New Delhi in collaboration with BIPL has launched "DrillBit : Plagiarism Detection Software for Academic Integrity" for the member institutions of DELNET. It is a sophisticated plagiarism detection software which is currently used by 700+ Institutions in India and outside. DrillBit is a global checker that uses the most advanced technology to catch the most sophisticated forms of plagiarism, plays a critical function for students and instructors and tag on a fully-automatic machine learning text- recognition system made for detecting, preventing and handling plagiarism and trusted by thousands of institutions across worldwide. DrillBit - Plagiarism Detection Software has been preferred for empanelment with AICTE and NEAT 3.0 (National Education Alliance for Technology) and contributing towards enhanced learning outcomes in India. On the other hand software uses a number of methods to detect AI-generated content, including, checking for repetitive phrases or sentences and AI-generated writing. As part of a larger global organization GJEIS (www.gjeis.com) and DrillBit better equipped to anticipate the foster an environment of academic integrity for educators and students around the globe. DrillBit is GDPR compliant with privacy by design and an uptime of 99.9% and have trust to be the partner in academic integrity (https://www.drillbitplagiarism.com) tool to check the originality and further affixed the similarity index which is {03%} in this case (See below Annexure 18.1.4). Thus, the reviewers and editors are of view to find it suitable to publish in this Volume 18, Issue-1, Jan-Mar 2026.

Annexure 18.1.4

Submission Date	Submission Id	Word Count	Character Count
21-Jan-2026	6050365 (DrillBit)	5494	38780

Analyzed Document	Submitter email	Submitted by	Similarity
3.1 CBS1_Sudhansh_ GJEIS Jan-Mar 2026. docx	sudhansh@ignou.ac.in	Sudhansh Sharma	03%

DrillBit

3	11	A
SIMILARITY %	MATCHED SOURCES	GRADE

A-Satisfactory (0-10%)
B-Upgrade (11-40%)
C-Poor (41-60%)
D-Unacceptable (61-100%)

LOCATION	MATCHED DOMAIN	%	SOURCE TYPE
1	files.eric.ed.gov	1	Publication
2	library.oapen.org	1	Publication
3	gurukuljournal.com	<1	Publication
4	files.eric.ed.gov	<1	Publication
5	athletesinaction.org	<1	Internet Data
6	eajournals.org	<1	Publication
7	repository.unika.ac.id	<1	Publication
8	Universities and Green Tech Leading the Transition to Sustainable Artificial I By Francisco J. Rey, Yr-2025,7	<1	Publication
9	meditation.mgh.harvard.edu	<1	Publication
10	research-advances.org	<1	Publication
11	www.everycrsreport.com	<1	Internet Data

Reviewers Memorandum

Reviewer's Comment 1: The manuscript addresses a highly relevant and emerging issue by examining the implications of Agentic AI for higher education institutions. The proposed distinction between operational resilience and pedagogical resilience provides a useful conceptual lens for understanding how universities can respond to the growing presence of autonomous AI systems. The paper is well-organized, supported by recent literature, and offers practical recommendations for institutional governance. As a minor suggestion, the authors may consider validating the proposed framework through empirical research or case studies in future work to strengthen its practical applicability.

Reviewer's Comment 2: The paper presents a timely discussion on the transition from generative AI to agentic AI and effectively highlights its implications for academic integrity, governance, and assessment practices. The dual resilience model proposed by the authors contributes to the ongoing discourse on responsible AI adoption in higher education. The recommendations for AI governance, critical AI literacy, and assessment redesign are particularly valuable for university administrators and educators. As a minor enhancement, the authors may consider including a conceptual figure illustrating the proposed dual resilience model to improve the clarity of the framework.

Reviewer's Comment 3: This manuscript makes a meaningful conceptual contribution to the emerging literature on Agentic AI in higher education by integrating governance and pedagogical perspectives within a single framework. The review of recent literature is comprehensive, and the discussion provides useful insights into the challenges and opportunities associated with autonomous AI systems in universities. The paper has practical relevance for policymakers, institutional leaders, and educators. For future research, the authors may consider examining the proposed framework across different institutional contexts and geographical regions to enhance its generalizability.

Citation

Sudhansh Sharma
"Navigating the Shift to the "Agentic AI University": Institutional Strategies for Operational and Pedagogical Resilience"
Volume-18, Issue 1, Jan-Mar 2026. (www.gjeis.com)

https://doi.org/10.18311/gjeis/2026
Volume-18, Issue 1, Jan-Mar 2026
Online iSSN : 0975-1432, **Print iSSN** : 0975-153X
Frequency : Quarterly, Published Since : 2009

Google Citations: Since 2009
H-Index = 96
i10-Index: 964

Source: https://scholar.google.co.in/citations?user=S47TtNkAAAAJ&hl=en

Conflict of Interest: Author of a Paper had no conflict neither financially nor academically.

Editorial Excerpt

The article has 03% plagiarism, which is the accepted percentage as per the norms and standards of the journal for publication. As per the editorial board's observations and blind reviewers' remarks, the paper had some minor revisions, which were communicated on a timely basis to the author (Sudhansh), and accordingly, all the corrections had been incorporated as and when directed and required to do so. The comments related to this manuscript are noticeably related to the theme "Navigating the Shift to the "Agentic AI University": Institutional Strategies for Operational and Pedagogical Resilience" both subject-wise and research-wise. The manuscript presents a timely and well-developed conceptual paper on the emerging role of Agentic AI in higher education. The proposed dual resilience framework offers a meaningful perspective for balancing AI-enabled innovation with human agency, academic integrity, and institutional governance. The study is supported by relevant and recent literature and provides practical recommendations that are valuable for university leaders, educators, and policymakers. After comprehensive reviews and the editorial board's remarks, the manuscript has been categorized and decided to be published under the "**Case-Based Study**" category.

Acknowledgement

The acknowledgment section is an essential part of all academic research papers. It provides appropriate recognition to all contributors for their hard work and effort taken while writing a paper. The data presented and analyzed in this paper by (Sudhansh) were collected firsthand, and wherever it has been taken, the proper acknowledgment and endorsement are depicted. The authors are highly indebted to others who facilitated accomplishing the research. Last but not least, endorse all reviewers and editors of GJEIS in publishing in the present issue.

Disclaimer

All views expressed in this paper are my/our own. Some of the content is taken from open-source websites, & some are copyright-free for the purpose of disseminating knowledge. Those are some we/I had mentioned above in the references section and acknowledged/cited as and where required. The author(s) have cited their own joint work mostly, and tables/data from other referenced sources in this particular paper with the narrative & endorsement have been presented within quotes and references at the bottom of the article. Accordingly & appropriately. Finally, some of the contents are taken from or overlap with open-source websites for knowledge purposes. Those are some of what I/we had mentioned above in the references section. On the other hand, opinions expressed in this paper are those of the author and do not reflect the views of the GJEIS. The authors have made every effort to ensure that the information in this paper is correct; any remaining errors and deficiencies are solely their responsibility