Excellence In Services International Conference EISIC
Formerly Toulon-Verona Conference, founded in 1998

Dealing with AI in Higher Education – How to go from theory to practice?

Anders Randler Uppsala University
Email: anders.randler@angstrom.uu.se
Raine Isaksson Uppsala University
Email: raine.isaksson@angstrom.uu.se

Abstract

Purpose: This paper examines the systemic integration of Artificial Intelligence (AI) within higher education, arguing that a proactive and transparent approach is essential for maintaining educational relevance, pedagogical value, and academic integrity.
Methodology: The study synthesizes current empirical research, meta-analyses, and policy frameworks regarding AI in higher education focusing on pedagogy with a detailed case study of practical, multi-layered implementations at Uppsala University, Sweden.
Findings: AI literacy courses for both students and faculty could effectively bridge the gap between theory and practice. Successful academical adaptation could be seen to rely on three critical pillars: (1) cultivating comprehensive AI literacy among educators and researchers to ensure pedagogical leadership, (2) integrating AI transparently across all curricula to align educational delivery with actual student practices, and (3) critically analysing how immediate, AI-generated proficiency affects students' cognitive development and deep learning.
Research Limitations/Implications: While the rapid evolution of generative AI necessitates continuously updated policy and pedagogy, early department-level interventions demonstrate that structured AI literacy courses for both students and faculty seems to effectively bridge the gap between theory and practice. Important to note is that this is a short and primarily conceptual study. The focus is to synthesise existing empirical research, meta-analyses, and policy guidance with a single departmental case study to develop a practical, transferable model.
Originality/Value: To move from theoretical anxiety to practical execution, this paper presents a concrete, scalable implementation model. This includes an analysis of Uppsala University's foundational student course, "Introduction to Generative AI and its Applications" (1TG334), a forthcoming advanced follow-up module, and the impact of professional development cycles for faculty The course "AI Literacy for Effective Teaching, Research, and Administration", offers a replicable roadmap for higher education departments.

1. Introduction

The public release and rapid adoption of widely accessible GenAI tools in late 2022 marked a major disruption in higher education. Rapidly, large language models (LLMs) and other generative technologies bypassed the traditional, deliberate channels of academic adoption, falling directly into the hands of students and faculty alike. Since then, the rapid pace of iteration—characterized by multimodal capabilities, autonomous agents, and highly specialized tutoring interfaces—has consistently outstripped departmental planning and regulatory adaptation (UNESCO, 2023).

This rapid technological evolution has placed universities in a state of pedagogical tension. On one hand, departments face the imperative to prepare students for a highly digitized workforce where AI literacy is increasingly demanded as a core competency (Almatrafi et al., 2024). On the other hand, the ease with which GenAI tools can produce human-like, technically proficient text, code, and creative media threatens the validity of traditional assessment models and introduces unprecedented opportunities for cognitive offloading (OECD, 2026). Early departmental responses frequently oscillated between extremes: from complete bans enforced by flawed automated detection software to a passive, unguided laissez-faire stance (Moore & Lookadoo, 2024).

However, historical precedents and emerging empirical evidence suggest that neither active resistance nor passive avoidance is sustainable; students are already using these tools on a massive scale to assist, and in some cases replace, their learning processes (OECD, 2026). Yet the existing knowledge remains fragmented. Much of it documents the capabilities and risks of GenAI in the abstract, or issues high-level policy guidance, while comparatively little work connects these strands into a coherent, actionable model that departments can implement and sustain. In particular, the literature has not adequately addressed how teachers and students can realistically adapt to a landscape that changes faster than department-level planning cycles, how AI can be integrated transparently without eroding the long-term relevance of education, or what the cognitive consequences of routine AI use are for students. It is this gap between abstract theory and resilient practice that the present study seeks to fill. Focus is on how existing course syllabi could be taught better using AI-support. The question of any changes required to the course content due to AI is not discussed.

To address this gap, this study is organized around three research questions (RQs) that together frame the transition from theory to practice:

RQ1: How do we in the academical context, as teachers and students, adapt to the new, rapidly changing academic landscape where AI is becoming increasingly central?

RQ2: How do we integrate AI in a transparent way that keeps our education relevant and future-proof?

RQ3: How does the use of AI affect students' cognitive development, and how should assessment adapt to protect it?

A note on scope is warranted. This is a deliberately short and primarily conceptual study rather than a large-scale empirical investigation. It does not report new experimental data or statistically test hypotheses; instead, it synthesises existing empirical research, meta-analyses, and policy guidance with a single-department case study to develop a practical, transferable model. This contribution should therefore be read as a structured roadmap for translating an emerging body of theory into resilient practice, and as a basis for subsequent empirical validation, rather than as definitive proof of the effect of any single intervention.

2. Literature Review and Theoretical Framework

2.1 Defining and Conceptualizing AI Literacy

In the context of modern higher education, literacy must extend beyond traditional reading, writing, and digital competencies to encompass "AI literacy." While early digital literacy focused on operating computers and navigating the internet, AI literacy demands an understanding of non-deterministic, probabilistic systems that interact, generate content, and make decisions autonomously.

In their systematic review of the literature from 2019 to 2023, Almatrafi et al. (2024) identify that AI literacy is not a single skill but a multidimensional construct. It spans technical dimensions (understanding machine learning, neural networks, and data biases), cognitive dimensions (the ability to critically evaluate, prompt, and collaborate with AI systems), and ethical dimensions (recognising issues related to intellectual property, data privacy, and societal impact). Similarly, Zhang et al. (2025) note in their umbrella review of systematic reviews that there is a growing consensus around defining AI literacy as a set of competencies that enable individuals to critically evaluate AI technologies, communicate and collaborate effectively with them, and make design-level decisions rather than remaining passive consumers.

A crucial finding in recent literacy research is that learning about AI must occur alongside learning with AI. Without a baseline understanding of how these systems function—including their limitations, structural biases, and "hallucination" tendencies—users are prone to over-reliance, automation bias, or outright dismissal of the technology (UNESCO, 2023). Building on this multidimensional view, the present study foregrounds what can be termed "AI cooperation": the collaborative dimension of AI literacy in which both student and teacher actively adapt their approach to learning and pedagogy through a commonly agreed partnership with AI. Understood this way, AI cooperation is not a substitute for AI literacy but one of its critical components, alongside the technical, cognitive, and ethical dimensions identified above. This study focuses on how such cooperation can be cultivated through the use of AI within existing university courses.

2.2 The Empirical Efficacy of GenAI in Education

To justify the integration of GenAI into curricula, it is necessary to examine its empirical impact on student learning results. Recent meta-analyses provide a nuanced picture of GenAI's educational efficacy.

In a comprehensive meta-analysis of generative AI interventions, Gu and Yan (2025) synthesized the effects of GenAI on student academic performance across both K-12 and higher education contexts. Analysing 19 empirical studies with 24 distinct effect sizes, they reported a highly significant overall effect size of Hedges, $g = 0.683$. This large effect size indicates that, on average, structured GenAI interventions significantly enhance students' academic achievements. Similarly, a meta-analysis by Liu et al. (2025) confirmed that GenAI applications can measurably boost learning results across diverse educational domains, particularly in terms of performance on standardized tasks, language learning, and technical problem-solving.

However, both meta-analyses highlight a critical caveat regarding how these gains are achieved. Gu and Yan (2025) performed a moderator analysis that yielded a striking contrast: students who engaged with GenAI with explicit teacher support and pedagogical scaffolding experienced exceptionally large performance gains ($g = 1.426$), whereas those who used GenAI without structured teacher support experienced negligible, statistically non-significant benefits ($g = 0.077$).

This empirical finding underscores that GenAI might not be a teaching cure-all that works on its own. Simply giving students access to AI models might not lead to genuine educational progress; rather, it could be the pedagogical design and instructor guidance surrounding the interaction that dictate whether the technology serves as a learning accelerator or a cognitive crutch.

2.3 Policy and Guidance in Higher Education

The rapid adoption of AI has forced departments to draft policies at an unprecedented rate. However, early policy attempts often focused on academic dishonesty and containment, relying on punitive language that alienated students and failed to account for the professional reality of AI integration.

Analysing these policy dynamics, Moore and Lookadoo (2024) advise that syllabus and department-level policies must pivot away from vague, prohibitive rules and toward transparent, audience-focused guidance. They argue that policy should be treated as a pedagogical tool: it must explicitly outline why and how AI can or cannot be used in specific assignments, aligning usage with specific learning objectives.

This aligns with the global guidelines published by UNESCO (2023) in its Guidance for Generative AI in Education and Research. UNESCO outlines a human-centred framework, asserting that the deployment of GenAI must safeguard human agency, intellectual diversity, and ethical validation. The guidance warns against the "homogenisation of opinions" and data privacy violations, advising educational departments to establish robust internal ethical review systems and clear age and content boundaries.

Furthermore, the OECD Digital Education Outlook 2026 highlights that the rise of GenAI demands a shift in the role of educational departments. The OECD argues that while AI can easily optimize task execution, it does not automatically guarantee "deep learning." Deep learning requires conceptual struggle, schema-building, and long-term memory integration. If a system completes the task for the student, the performance of the task is detached from the cognitive development of the learner. Thus, policy and curriculum must be designed to manage this trade-off, ensuring that AI enhances, rather than bypasses, the human cognitive effort required for learning.

2.4 Theoretical Framework

Two established frameworks anchor the analysis that follows. The first is the Technological Pedagogical Content Knowledge (TPACK) model (Yue et al., 2024), which holds that effective technology integration depends on the interaction of teachers' technological, pedagogical, and content knowledge rather than on the tool itself; it provides the lens for Pillar I's focus on educator competence. The second is the principle of "desirable difficulties" (Bjork & Bjork, 2011), which holds that the cognitive effort expended during learning is what makes knowledge durable and transferable; it provides the lens for Pillar III's concern with cognitive offloading and assessment. Read together, these frameworks connect the empirical evidence reviewed above to the three-pillar model developed in the Results, ensuring the practical recommendations rest on an explicit account of how people learn with technology.

3. Method

This study adopts a qualitative, case-informed synthesis design. It combines a structured review of empirical research, reviews, meta-analyses, and international policy guidance with an illustrative case from Uppsala University, Sweden. The case is situated within the University's emerging department-wide framework for generative AI, which includes university-wide guidelines, coordinated staff support, and an AI network for teachers. Within this broader framework, the study examines department-level initiatives developed at the Department of Civil and Industrial Engineering, Division of Quality Science, Campus Gotland. The principal local initiative is the five-credit first-cycle course *Introduction to Generative AI and its Applications* (1TG334), offered both at campus in Visby, Gotland and as a distance course. The campus course that was only open to program students had 35 students and the distance course had 102 students of which 55 actually took active part and 52 passed. The analysis also considers a planned advanced follow-up course and a professional-development initiative for staff "AI Literacy for Effective Teaching, Research, and Administration.". These initiatives are treated as local and emerging applications rather than as evidence of university-wide implementation. The case is used primarily to examine how teachers and students may adapt to AI and how transparent integration can be translated into educational practice. Evidence concerning cognitive development is derived principally from the research literature, while the case illustrates how identified cognitive risks and opportunities may be addressed through course design and assessment These cases were selected because, within a single department, they operationalise all three dimensions probed by the research questions: teacher and student adaptation (RQ1), transparent curricular integration (RQ2), and the cognitive effects of AI use (RQ3). The analysis maps evidence from both streams onto the three research questions and synthesises it into the multi-layered implementation model presented in the Discussion. Because generative AI is evolving rapidly, the model is treated as adaptive rather than fixed, with its components intended to be revisited as tools, policy, and empirical evidence continue to develop.

3.1 Tools and AI Use in the Research Process

In keeping with the paper's own argument for transparency, and in the spirit of "living as we learn," the AI tools used to produce this study are documented below. Four stages can be distinguished:

- **Research:** The initial search for relevant sources was carried out with Perplexity. The resulting articles were imported into Gemini Notebook to create a project that served as the basis for retrieval-augmented generation (RAG) in Gemini.
  To read and analyse the articles we used both the audio preview feature in Gemini Notebook and the app ElevenlabReader (reads a pdf as an audiobook).

- **Draft:** The first draft was produced in a Gemini canvas, with the Gemini Notebook project attached as context (RAG).
- **Writing process:** During writing and editing, Claude for O365 and ChatGPT Plus were used to generate ideas and alternative formulations.
- **Feedback and refining:** in the final stage, Claude for O365 was used again to adjust and improve the complete paper.

4. Results

This section presents the findings of the synthesis, organised around the three pillars that emerged as the foundation for proposed successful departmental adaptation. Each pillar addresses one of the study's research questions and combines evidence from the literature with observations drawn from the Uppsala University case studies: educator and researcher AI literacy (RQ1), transparent curricular integration (RQ2), and the cognitive effects of AI use together with the assessment reforms they motivate (RQ3).

4.1 Pillar I: Cultivating Teacher and Researcher AI Literacy

4.1.1 The Educator as the Instructional Gatekeeper

The empirical findings of Gu and Yan (2025) suggest that the educator is the most critical variable in the successful deployment of GenAI. If teachers do not understand the tool, they cannot provide the essential scaffolding required to elevate student performance.
Historically, professional development in educational technology has lagged behind technological advancement. With GenAI, this lag presents a severe risk. Faculty members who lack AI literacy may default to defensive pedagogical postures—such as reverting entirely to pen-and-paper exams without incorporating AI into their teaching, or ignoring student AI use altogether, leading to undetected academic dishonesty and unguided learning. To prevent this, departments must treat faculty upskilling as an immediate, strategic priority. Educators must develop "technological pedagogical content knowledge" (TPACK) that specifically incorporates AI (Yue et al., 2024), enabling them to design lessons where AI is used as a cognitive partner.

4.1.2 Case Study: Development Initiative from Uppsala University Campus Gotland

To address this critical need, the Division of Quality Science, Campus Gotland designed and executed a professional development course titled "AI Literacy for Effective Teaching, Research, and Administration." This initiative serves as a core framework for building departmental capacity, with the long-term goal of spreading the model to other higher education departments.

Course Objectives and Outcomes

The primary objective of the course is to equip faculty with the skills, conceptual understanding, and ethical frameworks necessary to integrate GenAI constructively into their daily work. The course defines four core targets:

1. **Prompt Engineering Mastery:** Training educators to transition from basic, single-turn queries to advanced, iterative prompt patterns (e.g., role-play prompting, chain-of-thought prompting, and contextual few-shot prompting) to maximize the quality and accuracy of AI outputs.
2. **Enhancing Distance and Campus Teaching:** Demonstrating how to use AI to build interactive, personalized, and engaging virtual classroom environments.
3. **Ethical and Technical Fluency:** Deepening faculty understanding of the underlying

technologies (such as transformers, parameters, and probabilistic generation) and the associated ethical challenges, including algorithmic bias, data harvesting, and intellectual property.

4. **Operational Efficiency:** Training teachers to utilize GenAI to streamline administrative burdens (e.g., draft communications, summarize meeting transcripts, synthesize reports) and support academic research (e.g., initial literature mapping, code generation for data analysis).

Structural Composition of the Educator/Researcher Course

From the beginning in the pilot versions of the course, it was organised in several smaller sessions. After feedback it was decided to create a curriculum organized into three central modules, balancing theoretical understanding with direct hands-on application. These should be given within a relatively short period of time. Each session is 3-5 hours:

Module Title	Focus Areas	Practical Deliverables
1. Generative AI in Education, Administration	Introduction with some historical context. Overview of major platforms (ChatGPT, Gemini, Claude, etc) and their use in daily work. Identifying hallucinations and biases.	Creation of a personalized "AI Tool Matrix" comparing capabilities across different commercial and open-source models.
2. Generative AI in Research and Development. Technical overview of LLM	Focusing on research tools as Perplexity, NotebookLM and Avidnote. Focus around how to get AI to use relevant info in research and development (using RAG). An overview of the technical basics of an LLM.	Creation of projects in different platforms in order to create reports/papers based on specific documentation and data avoiding hallucinations.
3. Ethical AI & Local AI	Introduction to core ethical issues in AI use — bias, privacy, academic integrity, and transparency. Overview of running AI locally (open-source/on-premise models) and when local AI is preferable to cloud-based tools for data privacy and control.	Assessing student work using local AI. Refining documents containing sensitive data using local AI.

4.2 Pillar II: Transparent Curriculum Integration

4.2.1 From Shadow Use to Explicit Integration

A major challenge currently facing higher education is the phenomenon of "shadow AI use"—where students rely on GenAI secretly to complete assignments, driven by a lack of clear guidance or fear of academic penalty. This hidden usage creates a disconnect between the student's actual learning workflow and the instructor's pedagogical design. Because the AI usage is hidden, the teacher cannot evaluate its impact or correct errors in prompt design, logic, or sourcing. To resolve this, universities must adopt a stance of transparency. If GenAI is a standard tool in the modern professional landscape, it must be treated as a standard tool in the academic landscape. This transition requires two synchronized actions:

1. formalising transparent AI policies in every course syllabus
2. introducing dedicated, credit-bearing courses that explicitly teach students how to use AI responsibly and effectively.

By integrating AI openly into the curriculum, educators can guide the student's learning journey directly. Rather than using AI to generate a final answer, students are taught to use it as a conversational partner, a brainstorming partner, a debugger, or an alternative perspective. This shifts the educational focus from the product of learning (e.g., the submitted essay or code) to the process of learning (e.g., the critical evaluation, editing, and validation of AI-generated content).

4.2.2 Case Study: "Introduction to Generative AI and its Applications" (1TG334)

The Department of Civil and Industrial Engineering (Uppsala University) has tested this approach by implementing a foundational undergraduate course, delivered both on-campus and via distance education: "Introduction to Generative AI and its Applications" (Course Code: 1TG334) (Uppsala University, 2025a). The campus course is only for program students but the distance course is open for all students. The course is structured to provide an immediate, accessible, yet academically rigorous entry point for students across all disciplines. As detailed in the official syllabus (Uppsala University, 2025b), the course centres on three primary intended learning outcomes (ILOs). Upon successful completion, students must be able to:

1. **Understand and Explain:** Account for basic terminology and the overall functionality, possibilities, and limitations of various types of artificial intelligence, with a specific focus on neural networks, machine learning, and training biases.
2. **Apply and Construct:** Use generative AI tools proficiently to create text, images, and other media based on advanced, multi-modal prompt patterns, including the synthesis of uploaded analytical materials.
3. **Reflect and Analyse:** Discuss diverse applications of generative AI across society and critically reflect on related ethical, social, legal, and copyright issues.

The pedagogical design of 1TG334 deliberately avoids requiring a background in computer science or programming. Instead, it utilises user-friendly interfaces to democratise access, ensuring that students from humanities, social sciences, and natural sciences alike can build technical competency. Teaching methods are highly interactive, utilising a blend of lectures, hands-on workshops, and collaborative seminars where ethical debates are foregrounded.

4.2.3 Case Study: The Next-Step Module: Continuous Adaptation

Recognizing that the AI landscape is in a constant state of flux, The Department of Civil and Industrial Engineering (Uppsala University) is currently developing a follow-up module building on the Introduction. This "next-step" course is designed to address the rapid obsolescence of specific software versions and interfaces by focusing on meta-skills and system adaptability.
This next-step curriculum focuses on:

- **Autonomous AI Agents:** Teaching students how to design, coordinate, and critique multi-agent workflows where separate AI agents perform specialized tasks (e.g., research, drafting, editing) under human supervision.
- **Domain-Specific AI Applications:** Transitioning from general-purpose LLMs to more specialised models in data science, creative writing, and legal analysis.
- **Retrieval-Augmented Generation (RAG):** Introducing students to the mechanics of feeding proprietary or highly specific academic datasets into local models to generate highly accurate, contextual analyses, bypassing the limitations of standard web training data.
- **General updates:** An overview of the current AI-tools and their updates referring back to the Introduction course

Through this two-tiered course structure, The Department of Civil and Industrial Engineering (Uppsala University) ensures that students do not just learn how to operate a specific version of a specific tool, but instead develop a durable, conceptual framework that allows them to adapt to whichever tools emerge in the future.

4.3 Pillar III: Cognitive Development, Assessment, and the Reintroduction of the Oral Exam

4.3.1 The Danger of Immediate Gratification: Cognitive Offloading vs. Deep Learning

While the technical and professional benefits of GenAI are clear, its impact on human cognition represents one of the most critical questions in contemporary educational science. Educational psychology has long established that learning is not a passive process of information acquisition; rather, it is an active construction of cognitive schemas that requires effort, attention, and cognitive struggle.
This concept, often described as "desirable difficulties" (Bjork & Bjork, 2011), suggests that the

cognitive labour expended during learning directly correlates with the durability of memory retention and the transferability of knowledge. When a student struggles to draft an essay, debug a line of code, or solve a mathematical equation, their brain is actively forming, testing, and modifying neural connections.

GenAI, by design, minimizes friction. It provides highly polished, technically accurate answers to complex questions in a fraction of a second. This introduces the profound risk of cognitive offloading—the systemic outsourcing of mental processing to external tools (OECD, 2026). If a student can bypass the struggle of drafting, reasoning, and error correction by simply asking a prompt, the immediate "task performance" may appear high, but the internal "learning outcome" could be severely compromised. On the other hand for the curious student AI provides ways of going further than what the educator might be capable of.

As the OECD Digital Education Outlook 2026 notes: "Better AI-supported task performance does not automatically mean deeper learning." If a student receives immediate gratification at every step of their inquiry, they might not develop the cognitive stamina, critical thinking, or deep conceptual understanding that defines higher education. Over time, this can lead to an erosion of independent analytical capability, leaving students highly proficient at generating content via AI, but incapable of evaluating its validity, explaining its logic, or generating original insights without technological assistance. However, restricting students from using AI as a discussion partner could slow down the learning of ambitious and talented students.

Through collaboration, GenAI interactions and learning tasks can be designed to preserve the productive struggle associated with “desirable difficulties” (Bjork & Bjork, 2011). Recent empirical evidence indicates that unrestricted AI assistance may improve immediate task completion while weakening knowledge retention or failing to produce corresponding learning gains (Barcaui, 2025; Graf et al., 2026). The central challenge is therefore not simply whether students use AI, but how teachers and students jointly structure its use so that it scaffolds rather than substitutes for cognitive effort (Bauer et al., 2025).

4.3.2 Rethinking Traditional Assessment Models

These cognitive concerns also raise questions about the continuing validity of traditional assessment practices. Higher education has long relied on out-of-class written assignments—including essays, take-home examinations, laboratory reports, and Bachelor’s, Master’s, and doctoral theses—as evidence of students’ knowledge and understanding. The increasing availability of GenAI does not necessarily make such assignments obsolete. However, it makes it more difficult to treat the submitted text alone as sufficient evidence of the student’s independent learning and cognitive development. GenAI can support activities ranging from language correction and feedback to generating arguments, analyses, and complete drafts. Consequently, an instructor examining only the final document may find it difficult to determine how the student produced the work, what role AI played, and to what extent the submission reflects the student’s own understanding. AI-detection tools offer only limited support in addressing this problem. Current detectors may produce false positives, may be less reliable for texts written by non-native English speakers, and can often be circumvented through editing or paraphrasing. Their results should therefore not be treated as definitive evidence of unauthorised AI use. A more constructive response is to reconsider how written work is designed, documented, and combined with other forms of assessment. This does not imply that traditional written assignments or supervised paper-based examinations should be abandoned. Guarded examinations can remain valuable when the purpose is to assess students' unaided knowledge and reasoning. Similarly, out-of-class written assignments remain important for assessing research, analysis, synthesis, and extended argumentation. The central issue is whether the final submitted text should stand alone as evidence of learning. Depending on the intended learning outcomes, it may be complemented by process documentation, transparent reporting of AI use, draft discussions, reflective commentary, presentations, or oral follow-up questions. The appropriate assessment design should therefore depend on what is being assessed and on whether, and for what purposes, the use of GenAI is permitted.

4.3.3 Combining AI-Supported Learning with Verified Assessment

The challenges associated with GenAI do not require one universal assessment solution. Depending on the intended learning outcomes, universities can combine several approaches, including supervised problem-solving, practical demonstrations, laboratory work, presentations, project defences, monitored written examinations, portfolio assessment, process documentation, and reflective accounts of how AI was used. Assessment should increasingly focus not only on the quality of the final product but also on the student's ability to explain, apply, evaluate, and defend the underlying knowledge and reasoning. Within this broader range of options, the present study gives particular attention to oral assessment. One possible approach is to distinguish more clearly between the learning phase and the final verification of learning. During the learning phase, students can be encouraged to use GenAI transparently as a learning resource. It may provide explanations, identify weaknesses in an argument, support literature exploration, generate alternative perspectives, or suggest problem-solving and debugging strategies. These benefits depend on students using AI critically and purposefully, supported by collaboration among teachers and students with different levels of AI experience. The objective should not be to remove all difficulty, but to use AI selectively while preserving the cognitive effort needed for durable learning. During the examination phase, the university may then use assessment methods that provide more direct evidence of individual understanding. Oral assessment is particularly promising for this purpose. Through follow-up questions, requests for clarification, and the application of knowledge to unfamiliar examples, the examiner can explore whether students understand the work they have submitted and can reason independently about its central concepts. Oral assessment may be used independently or in combination with a written assignment, project, practical task, or thesis. The proposal is therefore not to replace all existing examinations with oral examinations, nor necessarily to exclude AI from every assessment situation. Rather, oral assessment can provide a robust complement to AI-supported written work by enabling direct verification of student understanding. This combination preserves the benefits of GenAI during learning and knowledge production while strengthening confidence that the assessed learning outcomes have been achieved by the individual student.

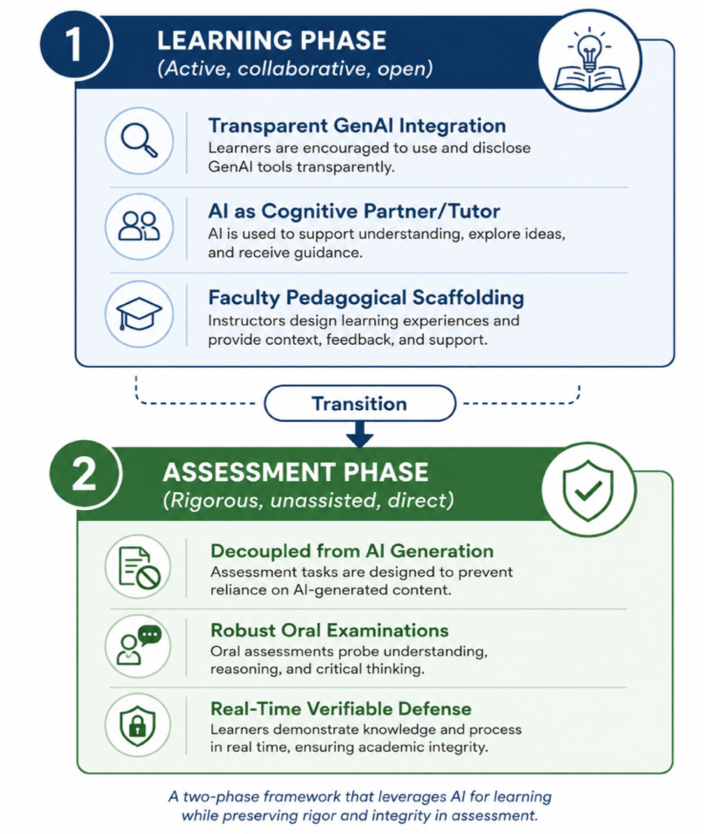

Figure 1. Decoupling the learning phase—where GenAI supports the student—from the examination phase, where an unassisted oral exam verifies independent understanding.

4.3.4 Oral Assessment as a Complement to AI-Supported Written Work

The expanded use of oral assessment at undergraduate, Bachelor's, and Master's levels represents one possible response to the difficulties of verifying individual learning in AI-supported written work. Oral examinations already have an established role in doctoral defences and in some professional programmes, but their wider applicability requires consideration of both their potential educational value and their practical limitations.

Oral assessment may contribute to several objectives:

1. **Direct exploration of understanding.** An oral examination enables the examiner to ask students to explain the methods, theoretical choices, arguments, and conclusions presented in their written work. Where GenAI has contributed to the production of a report or thesis, the discussion may provide additional evidence of whether the student understands and can independently account for its content. However, oral performance should not automatically be treated as a complete representation of subject knowledge, since students differ in their ability to communicate under examination conditions.
2. **Adaptive questioning.** Unlike a standardized written examination, an oral assessment allows questions to be adjusted in response to the student's answers. The examiner can request clarification, introduce alternative scenarios, and ask the student to apply concepts in

an unfamiliar context. This may help distinguish recall or reproduction from deeper conceptual understanding. At the same time, variations in questioning can reduce comparability between students unless the examination follows a sufficiently structured protocol.

3. **Assessment of communication and reasoning.** Oral examinations may assess competencies that are difficult to observe in a submitted document alone, including the ability to articulate an argument, respond to criticism, explain decisions, and reason in dialogue. These capabilities are relevant in many professional settings. Their inclusion in assessment should nevertheless be explicitly connected to the course learning outcomes rather than assumed to be universally relevant.
4. **Reduced dependence on the authorship of submitted text.** Synchronous oral interaction makes it more difficult for students to rely directly on externally generated answers. It is therefore relatively resilient to current forms of unauthorised GenAI assistance. It would, however, be too strong to describe oral examination as entirely immune to technological disruption. Students may prepare using AI, and future technologies could create new forms of assistance. The principal advantage is that the examiner can obtain additional, real-time evidence of the student's individual understanding.

4.3.5 Possible Models for Scaled Implementation

A major constraint on the wider adoption of oral examinations is the time and resources required, particularly in courses with large student cohorts. The following models illustrate how oral components could be incorporated without necessarily replacing existing forms of assessment.

- **Hybrid portfolio and oral assessment.** Students develop a written portfolio during the course and document how GenAI has been used. A structured oral component of approximately 15–20 minutes is then used to discuss selected elements of the portfolio. The oral component could contribute to, rather than necessarily determine, the final grade. Random selection of topics may encourage students to take responsibility for the portfolio as a whole.
- **Group presentation with individual questioning.** Groups of three or four students present a shared project, followed by questions directed to individual group members. This may reduce the time required while providing some evidence of individual understanding and contribution. Careful design is necessary to prevent group dynamics or unequal speaking opportunities from influencing individual grades.
- **Progressive development across a programme.** Oral assessment can be introduced gradually. Early courses might include short, low-stakes discussions or oral checks, while later courses could employ more formal project and thesis defences. Such progression may allow students to develop oral communication skills and become familiar with the format before it is used in higher-stakes assessment.

4.3.6 Limitations and Safeguards

The possible benefits of oral assessment need to be considered alongside significant concerns about validity, reliability, equity, and resources. Live oral examinations may disadvantage students experiencing performance anxiety, students with certain disabilities, and those working in a language in which they have less confidence. Outcomes may also be influenced by examiner expectations, differences in questioning, and unconscious bias. Furthermore, oral assessment can require substantial staff time.

Several design measures may reduce these risks:

- aligning questions and grading criteria explicitly with learning outcomes;
- using structured question frameworks and clear assessment rubrics;
- training and calibrating examiners;
- documenting the basis for grading decisions;
- providing appropriate accommodations and, where suitable, alternative formats;
- combining oral performance with written, practical, or process-based evidence;
- evaluating workload and assessment quality before wider implementation.

Oral assessment should therefore not be presented as a universal replacement for written assignments or supervised examinations. Rather, it can be investigated as one component of a diversified assessment system in which different methods provide complementary evidence of learning. In courses involving AI-supported written work, an oral component may strengthen the verification of individual understanding while also assessing the student’s capacity to explain and defend the work. The proposed approach is consequently best treated as a model for piloting and evaluation. Further empirical research is needed to determine its effects on assessment validity, student learning, equity, staff workload, and acceptance across disciplines. Its potential contribution lies in enabling universities to explore the benefits of transparent GenAI use during learning while retaining credible opportunities to assess students’ independently demonstrated knowledge and reasoning.

5. Discussion: The Multi-Layered Implementation Model

When we synthesize the three pillars analysed in this paper, a multi-layered model for systemic department-level adaptation emerges. Moving from theory to practice cannot be achieved through isolated, single-course interventions; it requires a coordinated ecosystem where student literacy, faculty competency, and robust assessment frameworks reinforce one another.

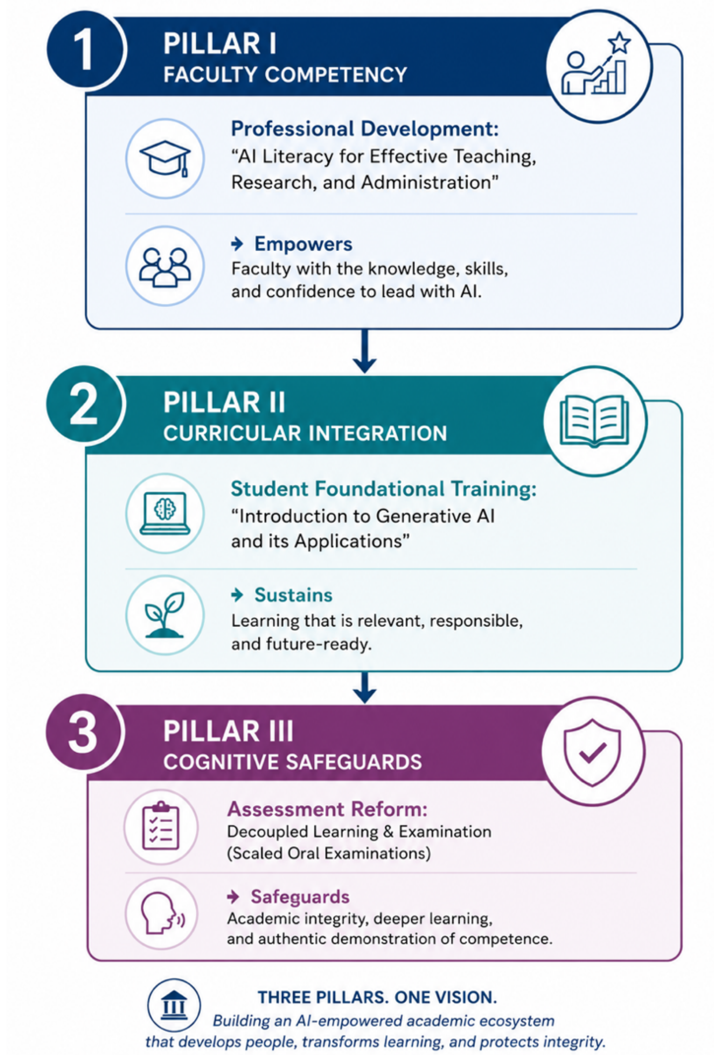

Figure 2. The multi-layered implementation model: an AI-literate faculty (Pillar I) enables transparent curriculum integration (Pillar II) and robust oral assessment (Pillar III), each reinforcing the others.

As illustrated in the model, Pillar I serves as the foundation. Without a highly trained, AI-literate faculty, it could be difficult to design, deliver, or update relevant student courses, let alone construct meaningful pedagogical scaffolding.

Once faculty competency is established, it flows directly into Pillar II, enabling the transparent, safe, and highly effective integration of GenAI across all curricula. Uppsala University's 1TG334 course provides the baseline student literacy, ensuring that students enter their major-specific courses with a shared, rigorous understanding of how to prompt, critique, and utilize AI ethically. This removes the "shadow use" dynamic, allowing educators to design coursework under the assumption that students are using AI tools openly and constructively.

Finally, Pillar III provides the essential cognitive and academic safeguards. It addresses the fundamental educational risk of GenAI: that immediate technological proficiency can bypass human cognitive growth. By decoupling the learning process (where AI is integrated as a supportive partner) from the examination process (where scaled oral exams verify independent understanding), universities can confidently maximize the benefits of the technology while ensuring that their graduates possess genuine, deep-seated intellectual capabilities.

This multi-layered model, as illustrated by the early experience at Department of Civil and Industrial Engineering (Uppsala University), proposes a resilient, scalable, and practical framework. It moves higher education away from the false dichotomy of banning versus blind acceptance, presenting instead a balanced path that uses technology to enhance human intelligence while safeguarding the core of human cognitive development.

These conclusions should be read with several limitations in mind. The account presented here is drawn primarily from a single department and reflects early, largely qualitative experience rather than controlled outcome data; the faculty and student initiatives described have not yet been evaluated against defined performance measures. The paper's central empirical claim—the contrast between scaffolded and unscaffolded GenAI use—rests heavily on a single meta-analysis with a limited number of underlying studies and small moderator subgroups, and should be treated as indicative rather than definitive. Finally, the model reflects the resourcing, regulatory context, and academic culture of the Swedish higher-education system; departments and faculties elsewhere may need to adapt it to different constraints. We therefore present this framework as a transferable starting point rather than a validated, universally applicable prescription.

6. Conclusion and Future Work

The integration of generative artificial intelligence in higher education represents a profound challenge to established academic structures, but it also offers an unprecedented opportunity to revitalise university pedagogy. As this paper has argued, successfully transitioning from theoretical concern to practical execution requires a systemic model built upon three interconnected pillars: upskilling educators, transparently integrating AI into the curriculum, and redesigning assessments to protect students' cognitive development.

Uppsala University's current initiatives – the "Introduction to Generative AI and its Applications" (1TG334) course for students and the "AI Literacy for Effective Teaching, Research, and Administration" program for faculty (developed and used by The Department of Civil and Industrial Engineering) – offer early, encouraging evidence that department-level change is achievable and can be scaled. When these educational offerings are combined with a shift toward direct, oral examinations, the threat of academic dishonesty and cognitive offloading is transformed into an opportunity for deeper human-to-human intellectual engagement.

The pilot cycles of the educator's course pointed to early, largely qualitative benefits. Faculty members reported not only a reduction in workload around basic administrative tasks but, more importantly, a marked increase in confidence when discussing AI use with students. Instead of policing AI, educators began using it to co-create course materials and design more relevant assignments.

The strategic roadmap involves scaling this professional development model. By refining the course material based on iterative faculty feedback, the department plans to offer the refined package consisting of three days to other departments. The plan for 26/27 is to offer the course locally at Campus Gotland to interested teachers and to two other departments in Uppsala (Mathematical Department and The Department of Ecology and Genetics). In the future the course will also be offered to other departments in Uppsala and for other interested organisations.

Looking to the future, the rapid progression toward autonomous, agentic AI systems will continue to challenge higher education. The traditional concept of a static degree program will likely give way to dynamic, lifelong learning frameworks where adaptability to emerging technologies is the primary skill. In this future landscape, the universities that thrive will not be those that built the tallest digital firewalls, but those that fostered deep, resilient AI literacy among their staff and students, whilst resolutely defending the value of the human struggle to learn, think, and understand.

References

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

Barcaui, A. (2025). ChatGPT as a cognitive crutch: Evidence from a randomized controlled trial on knowledge retention. *Social Sciences & Humanities Open*, *12*, p.102287.

Bauer, E., Greiff, S., Graesser, A.C., Scheiter, K. and Sailer, M. (2025). Looking beyond the hype: Understanding the effects of AI on learning. *Educational Psychology Review*, *37*(2), p.45.

Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In M. A. Gernsbacher, R. W. Pew, L. M. Hough, & J. R. Pomerantz (Eds.), *Psychology and the real world: Essays illustrating fundamental contributions to society* (pp. 56–64). Worth Publishers.

Graf, L., Bassner, P., Anzinger, M., Dietrich, F., Krusche, S. and Poquet, O. (2026). Interactive learning dashboards: rethinking learning visualisations as engagement tools. *Education and Information Technologies*, pp.1-27.

Gu, J., & Yan, Z. (2025). Effects of GenAI interventions on student academic performance: A meta-analysis. *Journal of Educational Computing Research*, *63*(1). https://doi.org/10.1177/07356331251349620

Gunsaldi, M. S., Guner, E. G., Uckan, M., & Bati, K. (2025). The impact of generative AI applications on student learning outcomes in science education: A systematic review. *Journal of Education in Science, Environment and Health*, *11*(3), 196–208. https://doi.org/10.55549/jeseh.840

Liu, X., Guo, B., He, W., & Hu, X. (2025). Effects of generative artificial intelligence on K–12 and higher education students' learning outcomes: A meta-analysis. *Journal of Educational Computing Research*, *63*(5). https://doi.org/10.1177/07356331251329185

Moore, S., & Lookadoo, K. (2024). Communicating clear guidance: Advice for generative AI policy development in higher education. *Business and Professional Communication Quarterly*, *87*(4), 535–556. https://doi.org/10.1177/23294906241254786

OECD. (2026). *OECD digital education outlook 2026*. OECD Publishing. https://www.oecd.org/en/publications/oecd-digital-education-outlook-2026_062a7394-en.html

UNESCO. (2023). *Guidance for generative AI in education and research*. UNESCO Publishing. https://doi.org/10.54675/EWZM9535

Uppsala University. (2025a). *Introduction to generative AI and its applications* (Course, Bachelor's level, 1TG334). https://www.uu.se/en/study/course?query=1TG334

Uppsala University. (2025b). *Introduction to generative AI and its applications: Syllabus, Bachelor's level, 1TG334*. https://www.uu.se/en/study/syllabus?query=51072

Yue, M., Jong, M. S.-Y., & Ng, D. T. K. (2024). Understanding K–12 teachers' technological pedagogical content knowledge readiness and attitudes toward artificial intelligence education. *Education and Information Technologies*, *29*(15), 19505–19536. https://doi.org/10.1007/s10639-024-12621-2

Zhang, S., Prasad, P. G., & Schroeder, N. L. (2025). Learning about AI: A systematic review of reviews on AI literacy. *Journal of Educational Computing Research*, *63*(3). https://doi.org/10.1177/07356331251342081