Scaffolding Educator AI Literacy: The SAGE Framework and the GUARD Learning Environment

Olga Rantanen[1] and Sebastian Moder[2]
[1]LAB University of Applied Sciences, Lappeenranta, Finland
[2]University of Liechtenstein, Vaduz, Liechtenstein

olga.rantanen@lab.fi
sebastian.moder@uni.li

Abstract: Generative artificial intelligence has moved from novelty to infrastructure in higher education within roughly two academic years, leaving educators caught between uncritical adoption and outright rejection. European policy frameworks increasingly call for AI-literate teaching, yet practical resources for translating high-level guidance into classroom practice remain scarce. This paper presents two outputs of the PATHFINDER project, an Erasmus+ Key Action 2 cooperation partnership involving the University of Liechtenstein, the University of Dubrovnik, LAB University of Applied Sciences, and formatio High School. The first output, Strategies for AI-Guided Education (SAGE), is an open instructional framework organised into four progressive modules: generative AI basics, prompt engineering, teaching with AI, and guiding students' AI use. The second output, the GUARD learning environment (Generative Universal Assistant for Resourceful Development), operationalises the framework through a hands-on web application that combines Socratic scaffolding with a live large language model chat interface. GUARD scaffolds effective AI use through five practice-oriented learning labs, each structured around three iterations focused on intent formulation, critical evaluation, and ownership. The first lab, Plan a Class Activity, supports the design of a learning activity for a specific topic, audience, and objective. The second, Create a Learning Assessment, supports the development of formative or summative items aligned to a target cognitive level. The third, Design a Rubric, supports the construction of analytic or holistic rubrics with specific, observable descriptors. The fourth, Revise an Existing Curriculum, allows users to upload a syllabus and revise it for stronger alignment, coherence, and AI literacy integration. The fifth, Revise an Assignment for AI Use, allows users to upload and redesign assignments so that student AI use is addressed fairly while preserving meaningful assessment. We describe the conceptual grounding, design rationale, content architecture, and technical implementation of both resources, and discuss their relevance to current debates on educator AI literacy, scaffolded skill acquisition, and responsible AI integration. The paper contributes a concrete, transferable response to a challenge faced by educational institutions worldwide.

Keywords: AI in Higher Education, AI Literacy, Educational Technology, DigCompEdu, Erasmus+

1. Introduction

In the span of two academic years, generative artificial intelligence has moved from a curiosity discussed in technology blogs to a piece of infrastructure that students now use routinely (Chan and Hu, 2023; Von Garrel and Mayer, 2023). The pace of this shift has caught many teaching faculty unprepared, leaving a sizable gap between the share of students who use generative AI tools regularly and the share of educators who feel confident teaching with or about them (Lee et al., 2024; Mah and Groß, 2024). The gap is not primarily a question of tool access, as chat-based AI is now free, widely available, and integrated into office software already licensed by many universities, but rather a question of pedagogical orientation. Educators are being asked to teach in an environment where a large language model can produce a passable first draft of an essay assignment in under a minute, yet few have been given a way to think through what this means for their courses, assessments, and relationships with students.

Two failure modes have become familiar in faculty common rooms (Lim et al., 2023; Zhai, 2025). The first is uncritical adoption, where educators accept AI outputs as authoritative or hand off pedagogically central tasks, such as drafting learning objectives, writing rubrics, or designing assessment items, to a model whose outputs they have not learned to evaluate. The second is blanket rejection, where educators ban AI use in their courses without offering students the literacy needed to engage with these tools responsibly elsewhere in their lives, while assessment design remains stuck in formats that the technology has rendered fragile. Neither response serves students well, and both leave educators in a defensive posture toward a technology that is unlikely to disappear.

The European policy response has been to articulate broad expectations. The DigComp 2.2 framework now incorporates AI-related examples of knowledge, skills, and attitudes within the digital competence framework for citizens (Vuorikari et al., 2022), and the DigCompEdu framework asks educators to develop and exercise digital pedagogical competence (Redecker, 2017). The European Commission's Ethics Guidelines for Trustworthy AI articulate principles such as human oversight, transparency, fairness, and accountability that educational

institutions are expected to uphold (European Commission, 2019). These documents are necessary, but they are not sufficient. They tell educators what to value; they do not tell them how to make those values operative when designing a lesson, drafting a rubric, or talking with a student about an assignment that may have been produced with AI assistance.

However, a persistent gap remains between policy-level guidance and day-to-day pedagogical action. Many educator resources take the form of institutional guidance documents, prompt collections, or short tool-demonstration workshops. These are useful starting points, but they often leave educators to make the crucial pedagogical translation themselves: how to formulate an AI-supported teaching task, how to evaluate the quality of an AI-generated output, how to recognise common failure modes, and how to adapt the result into a contextually appropriate classroom resource. The need is not only for information about AI but also for structured opportunities to practise the cognitive routines of responsible AI use.

The PATHFINDER project was conceived to address this gap. Funded under the Erasmus+ KA220-HED programme and running from January 2024 to December 2026, the project brings together partners from Liechtenstein, Croatia, and Finland to develop open, transferable resources for educator AI literacy across higher education and related educational settings.

Against this background, the paper is guided by three design questions. First, how can European AI literacy and digital competence frameworks be translated into practical educator-facing resources? Second, how can scaffolded interaction with large language models support educators' prompt formulation, output evaluation, and adaptation practices? Third, what design principles emerge from developing an open, transferable AI literacy learning environment for educators?

This paper reports on two of the project's main outputs: the SAGE framework (Strategies for AI-Guided Education), an open instructional framework for educators, and the GUARD learning environment (Generative Universal Assistant for Resourceful Development), a web-based application that operationalises the framework through scaffolded hands-on practice. This paper is positioned as a design and implementation report rather than an outcomes evaluation. Its purpose is to document the conceptual grounding, design rationale, content architecture, and technical implementation of SAGE and GUARD, and to identify design principles that may support adaptation by other educators and institutions. The paper contributes a transferable design account of how policy-level AI literacy principles can be translated into classroom-facing educator resources, with scaffolding treated as the central mechanism for developing responsible AI use.

2. Literature Review

The design of both SAGE and GUARD rests on three intersecting bodies of work.

The first is the European competence framework tradition. DigComp 2.2 (the digital competence framework for citizens; Vuorikari et al., 2022) and DigCompEdu (the digital competence framework for educators; Redecker, 2017) frame digital literacy as a layered set of competences that develop over time, from awareness and routine use to critical and creative engagement. DigCompEdu distinguishes between the digital competence educators need for their own professional practice and the competence they need to foster in learners. PATHFINDER's outputs were aligned with these frameworks from the start: SAGE's module structure reflects the DigCompEdu progression, while GUARD's iteration arc explicitly targets the higher levels of the DigComp 2.2 progression by asking users to evaluate, refine, and adapt AI outputs rather than merely consume them.

The second body of work is the EU's 2022 Ethics Guidelines for Trustworthy AI (European Commission, 2019), which set out seven requirements: human agency and oversight; technical robustness and safety; privacy and data governance; transparency; diversity, non-discrimination, and fairness; societal and environmental wellbeing; and accountability. Although these are normative principles, they can be translated into concrete design choices. GUARD's use of user-editable prompt templates, for example, directly reflects the principle of human agency: the scaffolding generates a prompt, but the user can always view, edit, and ultimately control what is sent to the model. Similarly, the decision to highlight failure modes such as sycophancy, hallucination, generic outputs, and vague descriptors reflects the principle of transparency by drawing attention to the ways AI outputs can fail in educational contexts.

The third body of work is scaffolding theory and its more recent extensions in metacognition and AI literacy research. Wood, Bruner, and Ross (1976) defined scaffolding as the support provided by a more capable other to enable a learner to accomplish a task they could not yet complete independently, with the expectation that this support is gradually withdrawn as competence develops. Vygotsky's (1978) zone of proximal development

provides the underlying developmental theory. Both ideas remain useful when applied to AI literacy: educators who use AI tools effectively are not those who memorise prompt tricks, but those who have internalised a way of thinking about prompts, outputs, and refinement. From this perspective, the role of scaffolding is to make that thinking visible long enough for it to become habitual. This is precisely GUARD's design wager. Mishra and Koehler's (2006) TPACK framework adds a complementary perspective: effective teaching with technology requires the integration of technological, pedagogical, and content knowledge, and weaknesses in any of these areas can lead to predictable failure modes. SAGE's later modules (Better Teaching with AI; Guiding Students' AI Use) are organised around these overlap zones, asking educators to consider AI integration in relation to their content domains and pedagogical commitments.

Recent AI literacy scholarship has converged on a multidimensional definition that goes beyond the ability to operate AI tools. Long and Magerko (2020) define AI literacy as a set of competencies that enable individuals to critically evaluate AI technologies, communicate and collaborate effectively with AI, and use AI as a tool. Subsequent work by Ng and colleagues (2021) distinguishes among the know-and-understand, use-and-apply, evaluate-and-create, and ethical dimensions of AI literacy. Educator AI literacy adds a teaching dimension to these: educators must not only be AI-literate themselves but also be able to develop AI literacy in their students. This dual demand is the explicit target of SAGE's fourth module (Guiding Students' AI Use).

A final strand of literature that shaped GUARD specifically is the empirical work on the failure modes of large language models in educational tasks. Recent literature documents several recurring patterns: sycophancy (the tendency of recent models to agree with user-suggested revisions even when those revisions reduce output quality; Sharma et al., 2024), hallucination (the production of plausible-sounding but factually inaccurate content, particularly in specialised domains; Ji et al., 2023), surface-level question generation (the tendency to default to recall-level assessment items even when higher-order thinking is requested), vague rubric language (descriptors that merely substitute adjectives across performance levels), and template-style activity generation (suggestions that sound engaging but could apply to almost any topic). GUARD's evaluation panels treat these failure modes as recurring phenomena that educators can learn to recognise. This stance—that model failure modes can themselves become a teachable curriculum—is one of GUARD's signature design moves.

3. Methodology

The development methodology was iterative. SAGE began with a needs analysis across the partner institutions, which identified a recurring tension between the pressure to engage with AI and the absence of structured guidance on what effective AI use should look like. The framework was reviewed against DigComp 2.2, DigCompEdu, and the EU AI ethics guidelines at each major iteration to ensure that its practical recommendations remained anchored in the European policy context.

GUARD followed a parallel but more software-intensive development process. Across both work packages, the deliberate design choice was to keep the resources lightweight, openly accessible, and easy to adapt. SAGE is published online as a structured set of information pages and worked examples that educators can read at their own pace and consult selectively when planning lessons. GUARD is deployed as a web application that requires no local installation and supports both a default model and a bring-your-own-key configuration. In the latter, educators or institutions can route requests through their own OpenRouter account and select a model that fits their data-governance requirements. Both resources are designed to remain usable beyond the funded period of the project.

4. Results

4.1 The SAGE Framework

SAGE's organising claim is that educator AI literacy can be developed by introducing the material in roughly the order in which educators are likely to need it. The framework is therefore arranged into four progressive modules, each containing a set of information pages and worked examples. The progression is linear in its default presentation but non-binding in use: experienced educators can enter at any module, and the framework is structured so that the modules cross-reference one another.

The first module, *Basics of (Generative) AI,* establishes the conceptual foundation. It opens with a plain-language explanation of generative AI, framed as a probabilistic next-token prediction system trained on large text corpora. This framing prepares educators to expect specific failure modes, such as hallucination, fluency without

correctness, or plausible filler, rather than be surprised by them. The module then turns to limitations and ethics, drawing on the EU AI ethics guidelines to discuss bias, accountability, and the limits of human oversight when an AI system produces outputs faster than a human can carefully evaluate them. A page on AI detection explains why current detection tools are unreliable in practice and why an assessment strategy that depends on detection without revising assignment design is fragile. The module closes with a page on when to use AI, which offers a working heuristic: AI assistance is most valuable for tasks in which the educator has the domain expertise to evaluate the output but lacks the time to produce it from scratch, and least valuable for tasks in which the educator lacks the expertise to recognise a flawed output. This heuristic recurs in different forms throughout the later modules.

The second module, *Prompt Engineering*, introduces the technical craft of communicating effectively with a language model. The pedagogical decision was to teach prompt engineering as a small set of transferable principles. The module covers five principles: assigning the model a job (which biases the model toward an appropriate register and stance); instructing the model throughout the prompt (which helps maintain coherence in longer prompts); using delimiters and special tokens to separate instructions from content (which reduces accidental instruction injection from pasted material); giving the model time to think (the practical implication of chain-of-thought prompting); and providing examples to guide the output (few-shot prompting, including examples of what to do and what to avoid). The module closes with a curated list of further resources for educators who want to explore the topic in greater depth, with the caveat that prompt engineering is a fast-moving area and that the underlying principles are more durable than any specific technique.

The third module, *Better Teaching with AI*, marks an inflection point in the framework. The first two modules describe the technology; the third asks educators to put it to work in service of pedagogical goals. The module is organised around five educator workflows: designing a curriculum, planning lessons, creating course materials, creating exam questions, and creating rubrics. Each workflow is presented through a worked example that pairs a realistic educator problem with a prompting and evaluation cycle. The pedagogical principle running through all five workflows is that AI output should not be treated as a finished product. In every example, the educator-author treats the AI's first response as raw material to be evaluated, refined, and adapted, while the worked example makes the evaluation criteria explicit. This creates the conceptual handoff to GUARD: the workflows SAGE describes in prose are operationalised as scaffolded labs in GUARD.

The fourth module, *Guiding Students' AI Use*, shifts the focus from educators' own use of AI to how they support students in using it effectively. The module is organised around five themes. Learning assessment discusses how to design assessments that remain meaningful when students have access to AI assistance, including the strategic use of process evidence, in-class oral defences, and assignments requiring local or personal context that the model cannot produce. Simulations describes how AI can be used to generate role-play partners, case scenarios, and counterfactual histories that would be expensive or impossible to produce manually. AI as a tutor discusses the use of AI as a one-on-one Socratic interlocutor for students working through difficult material, with attention to failure modes such as sycophancy and fabricated references. Evaluating AI outputs guides educators through a structured evaluation protocol that is appropriate to students' age and discipline. AI-assisted writing takes a developmental approach to the writing process, distinguishing between AI use that develops students' writing competence and AI use that merely outsources it.

Across all four modules, SAGE adopts a consistent voice: warm, substantive, opinionated where appropriate, and willing to acknowledge the limits of current knowledge. The framework treats AI as a powerful but unreliable tool whose responsible educational use depends on educators' pedagogical judgement, and it organises its content around developing that judgement.

4.2 The GUARD Learning Environment

GUARD takes the conceptual material that SAGE presents in prose and turns it into a hands-on practice environment. The key design decision is that the scaffolding is the product, while the LLM chat serves as supporting infrastructure. GUARD's value lies not in the chat interface itself but in the structured thinking it prompts before, during, and after each AI interaction. This section describes how that scaffolding is realised.

The application's core interaction model is a split-panel layout in which the left panel hosts the pedagogical scaffolding (reflection questions, a visible prompt template, evaluation tools, and a phase indicator), while the right panel hosts a chat interface with a large language model. The split creates a spatial separation between the cognitive work of preparing for and evaluating an AI interaction on the left and the act of conducting the interaction itself on the right. Educators using GUARD for the first time often comment that the layout slows

them down, which is precisely the intention. Whereas a default consumer-facing chat interface encourages rapid, unreflective interaction, GUARD's split-panel layout encourages a more deliberate workflow.

Within each lab, GUARD structures interaction around a three-iteration arc. The first iteration, *Get Something on the Page*, focuses on intent and clarity: the scaffolding helps the user articulate the pedagogical goal, student context, and constraints that the AI-generated output should respect. The second iteration, *Look Deeper*, focuses on critical evaluation and failure modes: the scaffolding shifts toward probing the AI output for problems and highlighting specific failure modes, including sycophancy, generic activities, vague descriptors, and hallucinated content, that recur in educational uses of generative models. The third iteration, *Make It Yours*, focuses on adaptation and ownership: the scaffolding guides the transition from AI output to usable classroom material by asking what the educator would change before using the output and what is missing that only the educator can supply. Across the three iterations, the scaffolding becomes progressively more demanding, reflecting a deliberate withdrawal of support as the user develops competence within the lab.

Within each iteration, GUARD guides the user through a four-phase sequence: *Reflect*, in which the user answers Socratic scaffolding questions whose responses populate a prompt template in real time in a mad-libs style; *Compose*, in which the user reviews and optionally edits the assembled prompt; *Generate*, in which the prompt is sent to the LLM and the response streams into the right panel; and *Evaluate*, in which the left panel displays a quality checklist, targeted reflection questions, failure-mode highlights, and, from the second iteration onward, a side-by-side comparison with the previous iteration's output.

The prompt template mechanism deserves particular attention because it embodies one of GUARD's core pedagogical principles. Each lab has a pre-defined template with labelled slots. In the activity-planning lab, for example, these include role and subject, student level, topic, learning objective, time constraint, and engagement type. As the user answers the scaffolding questions, the corresponding slots are populated in real time. The user can toggle between a structured view, in which the slots are visible and completed values are highlighted, and a raw view, in which the assembled prompt appears as a single editable block of text. The pedagogical aim is to make the link between structured thinking and effective prompting explicit: users repeatedly see, in real time, how their answers to Socratic questions map onto the structure of a well-formed prompt. The transferable skill is not simply the ability to use GUARD's templates, but the underlying habit of thinking in terms of role, audience, intent, constraints, and evaluation criteria before opening a chat interface.

The first version of GUARD includes five labs, each targeting a distinct educator workflow that recurs across the SAGE framework. *Plan a Class Activity* supports the design of a learning activity for a specific topic, audience, and objective. *Create a Learning Assessment* supports the development of quiz, exam, or formative assessment items with explicit attention to alignment and cognitive level. *Design a Rubric* supports the construction of an analytic or holistic rubric with specific, observable descriptors. *Revise an Existing Curriculum* supports the upload of a curriculum or syllabus and its revision for stronger alignment, coherence, and AI literacy integration. *Revise an Assignment for AI Use* supports the redesign of an existing assignment so that student AI use is addressed fairly while preserving meaningful learning assessment.

Each lab includes a distinctive set of failure-mode highlights that appear during the evaluation phase. The activity-planning lab warns about generic activities that sound engaging but amount to little more than a template with the topic name inserted. The assessment lab warns about the AI's tendency to default to recall-level questions even when higher-order thinking has been requested, as well as hallucinated content in specialised fields. The rubric lab warns about vague qualifiers that provide little guidance to students or graders and about symmetric descriptors that merely substitute adjectives across performance levels rather than describing qualitatively different kinds of work. These warnings reflect documented patterns in the empirical literature on LLM behaviour, and naming them gives educators a vocabulary they can carry into their own AI use outside GUARD.

5. Discussion and Conclusion

5.1 Discussion

The most important property of SAGE and GUARD is their complementarity. SAGE provides the conceptual map: the territory of teaching with AI, organised into modules and anchored in European competence frameworks and AI ethics guidelines. GUARD provides the practice ground: a structured environment in which this material can become habitual through scaffolded use. The two resources were designed in tandem and are best deployed

together, for example, in faculty development programmes that pair SAGE-based reading with GUARD lab sessions.

Treating scaffolding as the product, rather than the AI tool itself, has implications beyond PATHFINDER. Much current discourse about AI in education is organised around tools: which model, add-on, or detector is most reliable. This framing quickly becomes outdated and places the cognitive emphasis in the wrong place by encouraging educators to evaluate tools rather than outputs. GUARD's core design principle is that the durable skill is output evaluation and that making this work visible is more valuable than any specific tool integration. SAGE follows the same principle through tool-agnostic approaches to prompt formulation and output evaluation.

This framing also matters for faculty development. A one-off tool demonstration is poorly suited to a domain in which technology, policy, and pedagogical norms change quickly. PATHFINDER supports a different model: self-paced engagement with SAGE, structured practice in GUARD, and institutional support that allows the material to be revisited as technology evolves. This model is more demanding but also more robust because it builds educator capacity independently of specific tools or vendors.

Several broader debates intersect with the work reported here. First, SAGE and GUARD position critical capacity as a priority of AI literacy, with technical fluency viewed as necessary but insufficient and ethical awareness as a core component. Second, although both resources are educator-facing, they also develop educators' capacity to teach AI literacy to students; PATHFINDER's WP4, ARTISAN, will extend this work through a student-facing learning journey. Third, the GUARD assignment-revision lab and the SAGE module on guiding student AI use take the position that the question is no longer whether students will use AI, but how assignments and assessments can accommodate that use fairly while preserving learning and assessment validity.

5.2 Limitations and Future Work

Several limitations deserve acknowledgement. First, this paper documents the design and implementation of SAGE and GUARD but does not yet report evidence of their effects on educator practice or student outcomes. The next PATHFINDER phase, planned for autumn 2026, will evaluate SAGE and GUARD through face-to-face workshops for educators across the project partner institutions. Before each workshop, participants will receive selected SAGE materials for pre-reading, identify topics particularly relevant to their teaching practice, and prepare practical educational cases for use during the hands-on activities. Workshop content on SAGE will be adapted to participants' expressed interests, after which educators will work through the GUARD learning environment using authentic cases brought by participants or examples prepared by the workshop facilitators. Baseline and follow-up surveys administered before and after the workshops will be used to examine short-term changes in participants' confidence, perceptions, and other dimensions captured by the evaluation instrument. Quantitative pre/post data will be analysed to identify changes associated with the workshop intervention, while qualitative feedback and, where feasible, workshop artefacts will be used to identify perceived benefits, usability issues, and evidence of changes in prompt formulation, critical evaluation, and adaptation of AI-generated educational materials.

Second, the resources reflect a particular institutional and cultural context. PATHFINDER spans Liechtenstein, Croatia, and Finland, with a higher-education focus and one secondary education partner. The materials are in English; adaptation to other languages, disciplines, and educational systems is therefore a natural next step. Third, GUARD's five labs cover important educator workflows but do not exhaust them. Future labs are planned for feedback generation, differentiation, lecture preparation, and hallucination detection. A fuller implementation could explore adaptive scaffolding, in which evaluation questions respond to AI output rather than following a fixed schedule.

Finally, rapid technological change poses a maintenance challenge. GUARD's named failure modes are based on current models, and some may shift as models improve. Transferable failure-mode recognition partially mitigates this risk, but ongoing curation will remain necessary.

5.3 Conclusion

PATHFINDER set out to provide educators with a structured, transferable, and open response to a challenge that can no longer be treated as optional: integrating generative AI into teaching without resorting to either uncritical adoption or blanket rejection. SAGE and GUARD represent one such response. SAGE organises the conceptual territory, while GUARD provides a space in which educators can practise. Both are grounded in European

competence frameworks and AI ethics guidelines, designed for adoption and adaptation, and centred on scaffolding as a mechanism for responsible AI use. We present this work as a documented and transferable model.

Ethics Declaration

This paper reports the design and implementation of educator-facing project outputs and does not involve human-subject research, collection of personal data, or reporting of participant-level findings. Formal ethical approval was therefore not required.

AI Declaration

English is not the authors' first language; Grammarly was employed to ensure grammatical accuracy and improve stylistic consistency.

References

Chan, C.K.Y. and Hu, W. (2023) "Students' voices on generative AI: Perceptions, benefits, and challenges in higher education", *International Journal of Educational Technology in Higher Education*, Vol 20, No. 1, 43.

European Commission: Directorate-General for Communications Networks, Content and Technology and High-Level Expert Group on Artificial Intelligence (2019) *Ethics Guidelines for Trustworthy AI*, Publications Office of the European Union.

Ji, Z., Lee, N. and Frieske, R. et al (2023) "Survey of hallucination in natural language generation", *ACM Computing Surveys*, Vol 55, No. 12, pp 1-38.

Lee, D., Arnold, M. and Srivastava, A. et al (2024) "The impact of generative AI on higher education learning and teaching: A study of educators' perspectives", *Computers and Education: Artificial Intelligence*, Vol 6, 100221.

Lim, W.M., Gunasekara, A. and Pallant, J.L. et al (2023) "Generative AI and the future of education: Ragnarök or reformation? A paradoxical perspective from management educators", *The International Journal of Management Education*, Vol 21, No. 2, 100790.

Long, D. and Magerko, B. (2020) "What is AI literacy? Competencies and design considerations", *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp 1–16.

Mah, D.-K. and Groß, N. (2024) "Artificial intelligence in higher education: Exploring faculty use, self-efficacy, distinct profiles, and professional development needs", *International Journal of Educational Technology in Higher Education*, Vol 21, No. 1, 58.

Mishra, P. and Koehler, M.J. (2006) "Technological pedagogical content knowledge: A framework for teacher knowledge", *Teachers College Record*, Vol 108, No. 6, pp 1017–1054.

Ng, D.T.K., Leung, J.K.L. and Chu, S.K.W. et al (2021) "Conceptualizing AI literacy: An exploratory review", *Computers and Education: Artificial Intelligence*, Vol 2, 100041.

Redecker, C. (2017) *European Framework for the Digital Competence of Educators: DigCompEdu*, Punie, Y. (ed.), Publications Office of the European Union, Luxembourg.

Sharma, M., Tong, M. and Korbak, T. et al. (2024) "Towards understanding sycophancy in language models", *International Conference on Learning Representations*, pp 110–144.

Von Garrel, J. and Mayer, J. (2023) "Artificial Intelligence in studies—Use of ChatGPT and AI-based tools among students in Germany", *Humanities and Social Sciences Communications*, Vol 10, No. 1, 799.

Vuorikari, R., Kluzer, S. and Punie, Y. (2022) *DigComp 2.2: The Digital Competence Framework for Citizens - With New Examples of Knowledge, Skills and Attitudes*, Publications Office of the European Union.

Vygotsky, L.S. (1978) *Mind in Society: The Development of Higher Psychological Processes*, Harvard University Press, Cambridge, MA.

Wood, D., Bruner, J.S. and Ross, G. (1976) "The role of tutoring in problem solving", *Journal of Child Psychology and Psychiatry*, Vol 17, No. 2, pp 89–100.

Zhai, X. (2025) "Transforming Teachers' Roles and Agencies in the Era of Generative AI: Perceptions, Acceptance, Knowledge, and Practices", *Journal of Science Education and Technology*, Vol 34, No. 6, pp 1323–1333.