BUILDING AN AI TUTOR FOR UNIVERSITY LIFELONG LEARNING: DESIGN, IMPLEMENTATION, AND CRITICAL REFLECTION

Martin MAYR
BOKU University Vienna, Austria

Email: martin.mayr@boku.ac.at **Innovative Practice**

Keywords: *AI tutor; university lifelong learning; adult learners; personalised learning; innovative practice*

ABSTRACT

Artificial intelligence (AI) is increasingly promoted as a means to enhance flexibility, personalisation, and access in University Lifelong Learning (ULLL). This paper presents the development of an AI Tutor as an innovative practice designed to support adult learners in specialised university training, using Smart Farming modules as an illustrative context. Rather than positioning AI as a replacement for existing educational practices, the AI tutor was conceived as a pedagogically-bounded support layer embedded in established courses. Its design builds on curated course materials, lecture transcripts, and structured exercises, enabling context-aware, multilingual explanations and just-in-time learner support through a conversational interface. The paper outlines key design principles, implementation steps, and early practitioner observations through an iterative, design-based approach, highlighting perceived benefits such as increased accessibility, adaptability to heterogeneous prior knowledge, and integration of tacit instructional knowledge. At the same time, the contribution critically reflects on limitations and ethical considerations, including risks of over-reliance, authority attribution, data governance, and uneven learner readiness. The paper argues that the educational value of AI tutors in ULLL depends less on technical sophistication than on careful pedagogical curation, transparent design boundaries, and alignment with lifelong learning principles.

INTRODUCTION: AI AND UNIVERSITY LIFELONG LEARNING

Artificial intelligence (AI), specifically Large Language Models (LLMs), is increasingly positioned as a driver of change in University Lifelong Learning (ULLL). Technological transformation, labour market volatility, and demographic change have intensified demands for continuous upskilling across the lifespan. Universities are expected to provide flexible, scalable, and inclusive learning opportunities for diverse adult populations, and AI-enabled tools are often framed as a means to personalise pathways and extend institutional reach. However, effectiveness depends strongly on pedagogical design, implementation context, and ethical governance rather than on technology alone. From the perspective of adult learning theory, effective support must account for self-direction, experiential relevance, and heterogeneous prior knowledge, principles that any AI-based tool must actively accommodate (Reiss, 2021; Gupta and MacLellan, 2025).

Traditional digital formats in lifelong learning, notably Massive Open Online Courses (MOOCs) and static learning management systems, have achieved scale but frequently struggle to sustain engagement or address heterogeneous learner needs. Adult learners in ULLL often combine study with work and family commitments, requiring support that is available on demand and responsive to prior knowledge. Empirical evidence suggests that AI-driven tutoring systems can accelerate learning and provide timely feedback in both distance and adult education contexts, yet adoption remains uneven and often self-selective (Gupta and MacLellan, 2025; Möller et al., 2024). This raises a critical tension: while AI tutors show measurable benefits for motivated learners, they do not automatically resolve issues of access, inclusion, or persistence central to ULLL.

Recent research increasingly points toward hybrid and human-centred models of AI deployment rather than fully automated solutions. Studies on human–AI tutoring and co-pilot systems demonstrate that AI is most effective when it augments, rather than replaces, educators and trainers, supporting pedagogical quality while preserving human judgement and contextual sensitivity (Thomas et al., 2024; Wang et al., 2024). At the same time, systematic reviews highlight persistent risks related to algorithmic bias, transparency, and data governance, alongside a tendency for AI tools to focus narrowly on task performance rather than holistic self-regulated learning (Lan and Zhou, 2025). These findings underscore the need for critical, practice-oriented experimentation that aligns AI innovation with the values and goals of ULLL.

This paper presents the development of an AI Tutor as an innovative practice in specialised ULLL, offering a practice-oriented account aimed at instructional designers, ULLL programme managers, and university teachers seeking to integrate AI into existing course provision. Rather than presenting AI as a universal solution, it situates the AI Tutor within a specific educational context, examining its design rationale, practical implementation, and perceived added value for adult learners. By explicitly linking technological choices to ULLL principles, flexibility, inclusion, and learner-centeredness, the paper responds to current debates on how AI can meaningfully support lifelong learning without reproducing existing inequalities or pedagogical limitations. Its contribution is threefold: it applies AI tutoring to the under-explored context of adult university learners, proposes a reproducible workflow for repurposing existing course materials, and frames the AI Tutor as didactic middleware rather than a replacement for instruction.

CONTEXT AND PRACTICE SETTING

The practice described in this paper is situated within the Lifelong Learning unit of BOKU University Vienna, a university with a strong applied orientation towards sustainability, agriculture, and environmental sciences. The initiative was developed with partners from Technical University of Munich and Free University of Bozen-Bolzano in the context of the ERASMUS+ KA220 cooperation project "USAGE-NG" (an EU-funded multinational educational partnership; project reference 2022-1-AT01-KA220-HED-000089504; European Commission, 2022), which focuses on modernising university lifelong learning provision for non-traditional and regionally dispersed learners. Within this institutional setting, the AI Tutor was conceived not as a stand-alone technological experiment, but as an embedded innovation aligned with the university's mission to extend academic knowledge into professional and rural contexts.

The primary target group consists of adult learners with a university-level background, notably farmers and rural professionals engaged in continuing education in the domain of Smart Farming, covering topics such as precision agriculture, sensor technology, and data-driven crop management. These learners typically combine study with demanding work schedules and operate in environments characterised by limited access to on-campus support. Prior research on intelligent tutoring systems for adult learners suggests that flexibility, asynchronous access, and relevance to real-world practice are decisive factors for adoption and sustained use, particularly among older and professionally embedded learners (Gupta and MacLellan, 2025; Liu et al., 2025).

A key constraint shaping the design of the AI Tutor was the reliance on existing university courses and training materials. Rather than redesigning curricula from scratch, the innovation focused on restructuring and repurposing established Smart Farming modules into a conversational, context-aware format. This constraint reflects a common condition in ULLL, where scalability and sustainability depend on leveraging existing resources while enhancing personalisation and learner support through AI-driven tools (Möller et al., 2024).

DESIGN OF THE AI TUTOR

The design of the AI Tutor was guided by the dual requirement of pedagogical relevance for ULLL and practical feasibility within existing university course structures. Rather than conceiving the tutor as a generic conversational agent, the design process treated it as a form of didactic middleware: a dialogic interface layered onto established teaching materials that translates validated academic content into adaptive, context-sensitive explanations. This choice reflects a core assumption: in ULLL contexts, value lies less in novel content generation than in improving access, intelligibility, and adaptability of already validated academic resources.

Pedagogical design principles

Three pedagogical principles shaped the AI Tutor's design: learner-centredness, context sensitivity, and support for self-regulated learning. These principles draw on established adult learning theory, particularly the recognition that adult learners bring diverse prior knowledge, learn most effectively when content is relevant to their professional context, and benefit from opportunities for self-directed inquiry. First, learner-centredness was operationalised through adaptive explanations. The tutor was instructed to adjust the depth, terminology, and examples of its responses according to the learner's questions, thereby accommodating heterogeneous prior knowledge. This aligns with evidence that AI tutors can support personalisation and engagement in adult education settings, provided they are designed to respond flexibly rather than uniformly (Alfredo et al., 2024; Gupta and MacLellan, 2025).

The second design principle prioritised context sensitivity. The tutor does not operate on open web knowledge by default, but instead draws primarily on a curated, course-specific knowledge base derived from existing teaching materials. Technically, this is achieved through a Retrieval-Augmented Generation (RAG) architecture provided by the hosting platform (Academic AI, a GDPR-compliant cloud service for Austrian universities; ACOmarket, 2024), which retrieves relevant document passages before generating responses. This design decision significantly reduces, but does not eliminate, risks of hallucination and misalignment with course learning outcomes, while preserving coherence with institutional curricula. Importantly, the AI is positioned as a supplementary explainer rather than an authoritative instructor: a cautious stance consistent with evidence that AI-generated educational content requires human oversight and explicit quality controls (Pesovski et al., 2024).

Third, the AI Tutor was designed to support self-regulated learning by enabling iterative questioning, clarification, and reflection. Learners can revisit concepts at their own pace, request alternative explanations, or explore examples without the social cost sometimes associated with classroom questioning. While studies suggest that AI-supported self-regulation can enhance learning, they also warn against dependency risks if learners defer too much cognitive responsibility to the system (Lan and Zhou, 2025). This tension informed later design constraints.

Content preparation and structuring

Implementation began with a systematic collection of all available course materials. These included lecture slides, scripts, assignments, example exercises with solutions, past examination questions, and audio or video recordings of lectures. A key insight from this project is that AI tutoring quality is highly dependent on the coherence and organisation of source materials. Files were therefore renamed consistently and organised into a clear hierarchy by module and lecture unit, enabling traceability and future reuse.

To make the materials usable for an LLM, structured summaries were created for each lecture unit. Large Language Models were used to generate these summaries, with constraints to avoid excessive abstraction. Summaries focused on key concepts, definitions, and examples, often including references to the source material to preserve transparency and enable direct source attribution. This step reduces the context load for the AI system, allowing extended chats while maintaining response consistency and relevance.

Lecture recordings were transcribed using automated tools, followed by selective review. Transcripts served to capture what might be termed "tacit instructional knowledge," consisting of explanations, contextual examples, analogies, and clarifications that lecturers communicate orally but that are not recorded in slides or written materials. By transcribing and curating these elements, otherwise ephemeral knowledge becomes explicit and retrievable. Transcript-derived content was integrated into the summaries and clearly marked to maintain accountability for added interpretations.

Exercises formed a further implementation layer. Tasks were rewritten into clean, structured text files, including problem statements, step-by-step solutions, and grading hints where applicable. This enabled the AI Tutor to provide guided explanations, hints, or feedback rather than simply reproducing answers. The intention was to support self-regulated learning without automating assessment, a balance that is particularly important in adult and professional education contexts.

Defining tutor behaviour through a system prompt

Once the content was prepared, it was embedded into the AI system's knowledge base. The central design mechanism at the next stage was the system prompt, which defines the tutor's role, tone, and behavioural boundaries. The system prompt functions as a pedagogical instrument rather than a merely technical parameter. It instructs the LLM to respond in a supportive and explanatory manner, to adapt to different learner levels, and to rely primarily on the provided documents. *Table 1* below illustrates key behavioural rules encoded in the prompt.

Rule	Pedagogical intent
Respond based on provided course documents; use external knowledge only when necessary and flag it explicitly	Ensures curriculum alignment; reduces hallucination risk
Adapt explanations for beginner or advanced learners as appropriate	Accommodates heterogeneous prior knowledge
If unsure, say so	Promotes transparency and epistemic honesty
Highlight additional information derived from transcripts	Makes tacit instructional knowledge visible
Encourage reflection through follow-up questions rather than providing direct answers	Supports self-regulated learning

Table 1: Selected system prompt rules defining AI Tutor behaviour

Crucially, the prompt also directs the LLM to acknowledge uncertainty and to distinguish between information derived from the knowledge base and additional insights drawn from other sources or the model's training data. This design choice addresses ethical concerns related to overconfidence and perceived authority in AI systems, which can undermine learner judgement if left unchecked. By making the tutor's knowledge boundaries explicit, the implementation supports transparency and trust. For example, when a learner asks how soil moisture sensors function in precision agriculture, the tutor draws on module summaries and a transcript-derived analogy used by the lecturer, rather than generating a generic explanation from its training data (Dieterle et al., 2022).

Testing, refinement, and early use

The development followed an iterative, design-based approach. Rather than a single deployment, the AI Tutor underwent repeated cycles of testing, evaluation, and refinement by the project team (three instructional designers and subject-matter experts across the partner institutions). Testing consisted of structured interactions simulating typical learner queries like asking clarification questions, probing conceptual understanding, and requesting alternative explanations. Responses were assessed against a simple evaluation rubric focusing on clarity, factual accuracy, adaptability to learner level, and conversational coherence. Identified weaknesses led either to prompt refinements or to adjustments in how materials were summarised and structured.

Early use within the project context indicates that the AI Tutor is particularly effective as a just-in-time support tool. Learners can revisit specific concepts without navigating entire course documents and can engage with content in a more exploratory manner. These observations are consistent with broader evidence that AI-driven tutoring systems can improve efficiency and engagement in adult and distance learning when embedded within structured educational settings. However, it should be noted that these insights are based on practitioner observation during initial deployment rather than on systematic learner evaluation (Möller et al., 2024).

Critical reflection on implementation choices

A major strength of the described approach is its transferability. Because the workflow builds on existing teaching materials and institutionally approved AI platforms, it can be applied beyond the Smart Farming domain to any discipline where structured course materials, recordings, and exercises are available. The key requirements are not domain-specific expertise in AI, but rather pedagogical curation and willingness to restructure existing content for conversational use. This makes the approach particularly relevant for ULLL contexts, where programmes span diverse fields and resources for bespoke technological development are limited. At the same time, the approach is constrained by the quality of source materials; fragmented or outdated courses require substantial preparatory effort. Moreover, summarisation and transcription introduce interpretative layers that demand careful oversight to avoid subtle distortions.

Legal and ethical considerations further shape implementation. Some teaching materials can only be used within closed institutional systems due to copyright restrictions, limiting broader dissemination. Data protection and transparency requirements must also be addressed, particularly if learner interactions are logged or analysed. These constraints reinforce the importance of positioning the AI Tutor as a bounded support tool rather than an autonomous instructor.

BENEFITS, CHALLENGES AND LIMITATIONS

The perceived benefits of the AI Tutor emerged from its capacity to address persistent constraints in ULLL: heterogeneity of learners, time scarcity, and limited individualised support. The conversational interface enabled users to access explanations on demand, revisit concepts at their own pace, and pose questions without social barriers present in synchronous settings. A further added value lies in integrating formal course materials with tacit instructional knowledge from lecture transcripts, especially relevant in specialised domains such as Smart Farming, where learners benefit from concrete examples beyond static slides. The multilingual capabilities also supported inclusivity by lowering linguistic barriers. Rather than replacing educators, the tutor functioned as a complementary support layer, reinforcing the hybrid human–AI model that recent research suggests yields stronger outcomes than fully automated systems (Reiss, 2021; Gupta and MacLellan, 2025; Möller et al., 2024; Thomas et al., 2024).

These benefits must be interpreted cautiously. Observations are based on early-stage practitioner feedback rather than controlled experimental evaluation. There is a risk of overstating pedagogical impact if perceived improvements in engagement are conflated with demonstrable learning gains. Efficiency gains may reflect surface-level optimisation rather than deeper conceptual understanding, an ambiguity also noted in the wider literature (Ward et al., 2025).

Several challenges accompany implementation. Data protection and privacy remain central concerns, particularly given the handling of learner interactions. Compliance with GDPR and institutional data governance was treated as a non-negotiable design constraint. Content accuracy poses ongoing risks: while the tutor is grounded in curated materials, generative models can still produce plausible but incorrect explanations if safeguards are insufficient (An et al., 2024; Dieterle et al., 2022). Beyond technical constraints, the implementation rests on implicit pedagogical assumptions. Learners may attribute authoritative status to AI-generated responses regardless of explicit design boundaries. Continuous on-demand support, while intended to foster self-regulated learning, may equally reduce productive struggle if not carefully scaffolded. Multilingual availability does not automatically ensure inclusivity where disparities in digital confidence and AI literacy persist (Bulut et al., 2024; Meylani, 2025; Lan and Zhou, 2025).

Recognising these assumptions does not undermine the value of the AI Tutor; rather, it highlights the need for ongoing pedagogical mediation, explicit learner orientation, and institutional governance to ensure that AI-supported tutoring strengthens, rather than simplifies, ULLL practices.

CONCLUSIONS

The development of the AI Tutor illustrates how AI can extend learner support in ULLL without displacing established pedagogical practices. Rather than framing AI as disruptive, this contribution has shown how an AI Tutor can function as a pedagogically-bounded support layer that enhances access to established academic content. The educational value of such tools depends less on technical sophistication than on careful pedagogical curation, contextual embedding, and ethical governance (Dieterle et al., 2022; Reiss, 2021; Wang et al., 2024).

For ULLL, the implications are twofold. First, AI tutors can meaningfully extend institutional capacity to provide just-in-time, personalised support for adult learners who combine study with professional responsibilities. Observed benefits, increased accessibility, multilingual support, and integration of tacit knowledge align with evidence from adult and distance education contexts (Gupta and MacLellan, 2025; Möller et al., 2024). Second, these benefits are neither automatic nor universal. Without explicit pedagogical framing, learners may overestimate the authority of AI-generated explanations, and continuous support may unintentionally reduce productive struggle (Lan and Zhou, 2025). The implementation therefore supports a cautious, hybrid approach where AI tutors augment rather than replace human educators (Thomas et al., 2024; Wang et al., 2024).

Several research questions arise from this practice that warrant systematic investigation: How do adult learners perceive and navigate the knowledge boundaries of AI tutors in practice? What evaluation designs are appropriate for assessing AI tutor impact in ULLL settings where controlled experiments are difficult? Under what conditions does the described workflow transfer effectively to non-technical disciplines? Planned next steps include integrating guided exercises with scaffolded feedback, collecting structured learner feedback, and exploring more rigorous evaluation approaches. At this stage, the AI Tutor is best understood as infrastructure in the making: exhibiting properties such as persistence, scalability, and institutional embedding, but requiring further evaluation before it constitutes a mature component of lifelong learning ecosystems.

REFERENCES

ACOmarket (2024) *Academic AI.* Available at: https://acomarket.at/de/portfolio/projekte/academic-ai (Accessed: 24 March 2026).

Alfredo, R., Echeverria, V., Jin, Y., Yan, L., Swiecki, Z., Gašević, D. and Martinez-Maldonado, R. (2024) *Human-Centred Learning Analytics and AI in Education: a Systematic Literature Revie*', *Computers and Education: Artificial Intelligence*, 6. DOI: 10.1016/j.caeai.2024.100215

An, Q., Yang, J., Xu, X., Zhang, Y. and Zhang, H. (2024) 'Decoding AI ethics from Users' lens in education: A systematic review', *Heliyon*, 10(20). DOI: 10.1016/j.heliyon.2024.e39357

Bulut, O., Beiting-Parrish, M., Casabianca, J. M., Slater, S. C., Jiao, H., Song, D., Ormerod, C. M., Fabiyi, D. G., Ivan, R., Walsh, C., Rios, O., Wilson, J., Yildirim-Erbasli, S. N., Wongvorachan, T., Liu, J. X., Tan, B. and Morilova, P. (2024) *The Rise of Artificial Intelligence in Educational Measurement: Opportunities and Ethical Challenges*, Chinese/English Journal of Educational Measurement and Evaluation, 5(3). DOI: 10.59863/MIQL7785

Dieterle, E., Dede, C. and Walker, M. (2022) *The cyclical ethical effects of using artificial intelligence in education*, AI and Society, pp. 1–11. DOI: 10.1007/s00146-022-01497-w

European Commission (2022) *USAGE-NG — ERASMUS+ Project.* Project reference: 2022-1-AT01-KA220-HED-000089504. Available at: https://erasmus-plus.ec.europa.eu/projects/search/details/2022-1-AT01-KA220-HED-000089504 (Accessed: 24 March 2026)

Gupta, A. and MacLellan, C. (2025) *Intelligent Tutors Beyond K-12: An Observational Study of Adult Learner Engagement and Academic Impact*. arXiv preprint. DOI: 10.48550/arXiv.2502.16613

Lan, M. and Zhou, X. (2025) *A qualitative systematic review on AI empowered self-regulated learning in higher education*, npj Science of Learning, 10(1), p. 21. DOI: 10.1038/s41539-025-00319-0

Liu, V., Latif, E. and Zhai, X. (2025) *Advancing Education through Tutoring Systems: A Systematic Literature Review*. arXiv preprint. DOI: 10.48550/arXiv.2503.09748

Meylani, R. (2025) *Rethinking explainability in educational artificial intelligence: a critical systematic review of models, applications, and ethical dimensions*, Veredas do Direito, 22(3), p. e223553. DOI: 10.18623/rvd.v22.n3.3553

Möller, M., Nirmal, G., Fabietti, D., Stierstorfer, Q., Zakhvatkin, M., Sommerfeld, H. and Schütt, S. (2024) *Revolutionising Distance Learning: A Comparative Study of Learning Progress with AI-Driven Tutoring*. arXiv preprint. DOI: 10.48550/arXiv.2403.14642

Pesovski, I., Santos, R., Henriques, R. and Trajkovik, V. (2024) *Generative AI for Customizable Learning Experiences*, Sustainability, 16(7), p. 3034. DOI: 10.3390/su16073034

Reiss, M. J. (2021) *The use of AI in education: Practicalities and ethical considerations*, London Review of Education, 19(1). doi: 10.14324/LRE.19.1.05

Thomas, D. R., Lin, J., Gatz, E., Gurung, A., Gupta, S., Norberg, K., Fancsali, S. E., Aleven, V., Branstetter, L., Brunskill, E. and Koedinger, K. R. (2024) *Improving Student Learning with Hybrid Human-AI Tutoring: A Three-Study Quasi-Experimental Investigation*, in Proceedings of the 14th Learning Analytics and Knowledge Conference. New York, NY: ACM, pp. 404–415. DOI: 10.1145/3636555.3636896

Wang, R. E., Ribeiro, A. T., Robinson, C. D., Loeb, S. and Demszky, D. (2024) *Tutor CoPilot: A Human-AI Approach for Scaling Real-Time Expertise*. arXiv preprint. DOI: 10.48550/arXiv.2410.03017

Ward, B., Bhati, D., Neha, F. and Guercio, A. (2025) *Analyzing the Impact of AI Tools on Student Study Habits and Academic Performance*, in 2025 IEEE 15th Annual Computing and Communication Workshop and Conference (CCWC). pp. 434–440. DOI: 10.1109/CCWC62904.2025.10903692