



# USAGE-NG

Up-skilling Agricultural Engineering  
Next Generation

## AI Tutor Core Manual

Building an AI Tutor for University Courses: A Practical  
Workflow

Work Package n°: **5**

Activity n°: **5.3**

Date - Version: **11-2025 – 1.6**

Document prepared within the USAGE-NG Project with the support of the Erasmus+ Programme of the European Union: 2022-1-AT01-KA220-HED-000089504.

The European Commission's support for the production of this publication does not constitute an endorsement of the contents, which reflect the views only of the authors, and the Commission cannot be held responsible for any use which may be made of the information contained therein.

## Contents

1	Summary.....	3
2	Rationale behind an AI Tutor.....	4
3	How to Prepare a Course for AI Tutor Integration.....	4
3.1	Collect All Course Materials.....	4
3.2	Summarize Materials.....	5
3.3	Transcribe Recordings.....	5
3.4	Integrate Transcript Content.....	5
3.5	Prepare Exercises.....	6
3.6	Build the AI Tutor.....	6
3.7	Test and Refine.....	7
4	Copyright Considerations.....	7
5	Choosing a LLM.....	7
6	Checklist for Educators.....	8
7	System Prompt Example.....	8
8	System Prompt Design and Pedagogical Coherence.....	10

## 1 Summary

Within the USAGE-NG project, a structured workflow was developed to prepare existing university teaching materials so that a Large Language Model (LLM), such as ChatGPT, Mistral, or a locally hosted model, can function as a digital tutor for course content. The workflow focuses on pedagogical preparation rather than technical complexity and includes collecting all relevant materials, creating structured summaries, transcribing lecture recordings, and organising exercises in a consistent and transparent way. These materials are embedded into a customised Artificial Intelligence (AI) system, resulting in a chatbot that can explain concepts, answer questions, and respond flexibly to different learner needs based on curated course content.

The AI Tutor is designed to act as a supportive explanatory layer rather than an authoritative instructor. By grounding responses in validated academic materials and clearly defined system prompts, the tutor can adapt explanations to learners' prior knowledge, provide just-in-time clarification, and integrate insights from lecture transcripts that are often absent from static slides. This dialogical approach addresses common challenges in University Lifelong Learning, such as heterogeneous learner backgrounds, limited time availability, and the need for flexible, asynchronous support, while preserving alignment with institutional curricula and learning objectives.

Importantly, the workflow emphasises responsible and sustainable use of AI in education. Educational value is understood to arise not from automation alone, but from careful pedagogical design, transparent use of sources, and clear boundaries between human and artificial roles. Legal, ethical, and organisational considerations, such as copyright restrictions, data protection, and institutional governance, are treated as integral to the implementation process. Positioned in this way, the AI Tutor complements existing teaching practices and supports core principles of University Lifelong Learning: learner-centredness, flexibility, inclusivity, and long-term sustainability within established university structures.

## 2 Rationale behind an AI Tutor

University Lifelong Learning is characterised by heterogeneous learner profiles, limited time availability, and a strong need for flexible, just-in-time support. While digital formats such as learning management systems or MOOCs have increased access, they often remain static and poorly adapted to individual questions, prior knowledge, or situational constraints of adult learners. An AI Tutor addresses this gap by shifting the focus from content delivery to dialogue: learners can ask targeted questions, request clarification, and revisit concepts at their own pace. In this sense, the AI Tutor functions as a continuously available explanatory layer that enhances accessibility and persistence without requiring learners to navigate entire courses for specific needs.

The rationale for using an AI Tutor is deliberately cautious rather than techno-optimistic. Educational value does not arise from automation alone, but from careful pedagogical design, transparent use of curated content, and clear boundaries between human and artificial roles. Positioned as a support tool rather than an authority, the AI Tutor augments existing teaching practices by making validated academic knowledge more accessible and adaptable, while preserving human oversight and ethical responsibility. This approach aligns AI use with core lifelong learning principles: flexibility, learner-centredness, and sustainability within established university structures.

## 3 How to Prepare a Course for AI Tutor Integration

Preparing a course for AI Tutor integration is primarily a pedagogical and organisational task, not a technical one. The quality of an AI Tutor depends less on the underlying model than on how clearly, consistently, and transparently course materials are structured and contextualised. This section outlines a practical workflow for transforming existing teaching resources into a coherent knowledge base that enables the AI Tutor to explain content accurately, adapt to learner needs, and operate within clearly defined pedagogical and ethical boundaries.

### 3.1 Collect All Course Materials

Gather all available resources including:

- Lecture slides (PDF, PowerPoint)
- Scripts, manuals, and handouts
- Assignments and example exercises (with solutions)
- Past exam questions
- Audio and video recordings of lectures

Organise and name files systematically for future retrieval and reuse by establishing a clear hierarchy and consistent naming conventions, such as including course codes, topic identifiers, and version numbers. Store documents in folders labelled by module or lecture, and use metadata fields or simple tagging systems to describe content, author, and date. This structure not only simplifies AI integration but also supports long-term curation and sharing within academic teams.

### 3.2 Summarize Materials

Use ChatGPT or another LLM to create structured summaries for each unit. Focus on clarity and include headings, bullet points, and key terms. Example prompts:

*“Please summarize this lecture content with structured formatting and clear subheadings. Structure along slide numbers and do not summarize more than 2 slides in a section.”*

*“Summarize the lecture content into concise, clearly labelled sections, following slide order. Use bullet points for key ideas, definitions, and examples. Include 2–3 key terms per section and maintain consistent terminology with the source material. Clearly reference the slide number where each point originates to maintain traceability. Aim for balance—capture essential context without losing clarity or brevity.”*

Summarisation helps AI systems handle large volumes of text efficiently and improves coherence by distilling essential information. A good prompt for summarisation should specify both scope and focus. For instance, instruct the AI to highlight central concepts, identify examples, define technical terms, and note connections between ideas. Example additions could include: “Include the slide number for each point,” “List two key terms per section,” or “Emphasise conceptual relationships and omit repetitive phrasing.” These details guide the model to produce summaries that are contextually faithful, pedagogically useful, and ready for integration into AI tutoring workflows.

### 3.3 Transcribe Recordings

Use transcription tools such as noScribe (Open Source, <https://github.com/kaixxx/noScribe>) or commercial alternatives for high-quality, multilingual transcription. Transcripts should be clearly named and ordered according to the lecture sequence. In addition to basic organisation, pay attention to transcription accuracy, language support, and ethical considerations. High-quality transcriptions should preserve technical terminology and avoid misinterpretations caused by accents or noise. When possible, use human review to correct automated outputs and ensure multilingual accessibility. Always secure consent from speakers before recording, anonymise sensitive data, and comply with institutional data privacy policies, such as GDPR, to maintain ethical and legal integrity.

### 3.4 Integrate Transcript Content

Enrich the summaries from step 2 with insights from corresponding transcripts. Use a prompt like:

*“Please enhance the following summary with relevant information from the transcript. Highlight all new content from the transcript in bold and italics.”*

*“Integrate relevant information from the transcript into the summary while maintaining the original structure. Add time stamps or slide references for each new detail to ensure traceability. Prioritise explanations, examples, or clarifications that expand understanding. Clearly distinguish new insights using bold and italics, and merge overlapping information smoothly without duplicating content.”*

Including transcript-derived details can bring the AI tutor closer to the authentic classroom experience. For instance, a transcript might capture a lecturer’s spontaneous clarification of a complex concept, or an anecdote connecting theory to real-world

application. These elements enrich summaries by preserving nuance, tone, and contextual cues that slides alone cannot convey. Incorporating such insights helps learners engage more deeply and supports adaptive, conversational tutoring.

### 3.5 Prepare Exercises

Document tasks and exercises in clean, structured text files. Include:

- Problem statement
- Step-by-step solution or model answers
- Grading hints or rubrics (if applicable)

Exercises can be transformed into interactive formats that respond dynamically to learner input. For instance, AI tutors can provide progressive hints rather than full answers, offer automated feedback on multiple-choice questions, or adjust problem complexity based on learner performance. Embedding conversational prompts such as ‘Would you like a hint or a full explanation?’ encourages reflection and self-regulation. Additionally, integrating scoring rubrics and feedback templates allows learners to understand evaluation criteria in real time, turning static exercises into adaptive learning experiences.

### 3.6 Build the AI Tutor

Refer to the chapter on *Choosing an AI or LLM* at the end of this guide for detailed guidance on selecting an appropriate model and platform.

Upload all materials created in the previous steps to the LLM’s knowledge base.

Define the tutor’s role and behavior clearly. This is called a “System Prompt.” Example:

*“You are a tutor for the course ‘Introduction to Remote Sensing’. Respond in a supportive and explanatory tone, and adapt explanations for beginner or advanced learners as appropriate. You answer based on the provided documents. Use other information only if necessary. If unsure, say so. Highlight extra information from transcripts. Adjust answers to the learner’s level.”* (A longer example is in the demo files.)

A well-designed system prompt can also briefly specify parameters like the tutor’s communication style, level of formality, and feedback approach. For instance, it might instruct the AI to ‘use examples relevant to adult learners,’ ‘encourage critical thinking through short follow-up questions,’ or ‘summarise key points at the end of each response.’ Such details keep the prompt concise while ensuring that the AI behaves consistently in line with pedagogical goals. The prompt should also define clear boundaries, for example, topics the AI must avoid or defer to a human lecturer on—and may include a short list of covered topics or learning units. It can also reference external resources such as lecture slides, readings, or demo files to support accurate and contextualised answers.

To ensure pedagogical coherence, this section could discuss how AI tutors support structured learning objectives by mapping their responses to defined learning outcomes. For instance, prompts could encourage the AI to use Bloom’s taxonomy verbs such as ‘explain,’ ‘apply,’ or ‘evaluate’ to guide learners toward deeper understanding. Furthermore, linking the design to University Lifelong Learning (ULLL) principles would

highlight adaptability, learner autonomy, and reflection as integral parts of the AI's behavior. Providing short examples of how a system prompt aligns with each cognitive level would make this design approach concrete and replicable.

A short example of a minimal System Prompt is at the end of this document.

### **3.7 Test and Refine**

Interact with the AI as a learner would. Ask content questions and check the AI's ability to:

- Explain clearly and appropriately
- Handle follow-up questions
- Adapt explanations to different levels of prior knowledge
- Provide feedback on exercises (if included)

Creating a simple evaluation rubric ensures a consistent and objective way to test AI tutor quality. For example, evaluators could assign scores (1–5) to dimensions such as factual accuracy, clarity of tone, adaptability to learner level, and engagement in dialogue. Including brief descriptors (e.g., '5 = consistently accurate and contextually aware; 1 = frequent factual errors') makes the assessment transparent and repeatable. This framework allows for both quantitative and qualitative evaluation and helps guide iterative improvements of the AI tutor's design.

## **4 Copyright Considerations**

When preparing materials for integration into an AI tutor, it is essential to distinguish between content that may be used internally for teaching and content that can be shared or published. Many teaching materials (e.g., slides from publishers, licensed journal articles, or third-party images) are legally usable within university teaching environments but must not be made publicly accessible. This includes uploading to publicly shared GPTs, websites, or open-source platforms. In such cases, materials should only be embedded into closed systems (e.g., institutional Tailored AIs) with restricted access for enrolled students. Whenever possible, use openly licensed or self-produced content. If in doubt, consult your institution's copyright office before uploading any materials.

## **5 Choosing a LLM**

Selecting an appropriate Large Language Model (LLM) or AI platform is a crucial step that should consider both pedagogical aims and institutional constraints. The choice depends on factors such as privacy requirements, cost, technical support, and integration capabilities. Many universities have policies that specify which AI tools can be used for teaching; educators should consult their IT or e-learning departments to confirm approved options. When institutional approval is required, documentation about data handling, API usage, and model hosting should be reviewed carefully to ensure compliance with university and national regulations.

Practical factors to consider include:

- **Data privacy and storage:** Verify where data is processed and stored, especially if student data is involved.
- **IT support:** Confirm the level of technical help available for deployment, troubleshooting, and maintenance.
- **Ease of use:** Evaluate how intuitive the platform is for non-technical educators.
- **Customization:** Check whether prompts, access permissions, and datasets can be tailored for course-specific needs.
- **Costs and licensing:** Clarify subscription fees or token usage policies early on.

When uncertain, start with institutionally supported platforms (e.g., AcademicAI or Microsoft Copilot for Education, ChatGPT for Education) before experimenting with independent models like Mistral, ChatGPT or Claude. This reduces compliance risk while allowing gradual adoption of AI-enhanced tutoring.

## 6 Checklist for Educators

- Verify that all course materials are either self-produced or carry an open licence (e.g., Creative Commons).
- Confirm usage rights for third-party images, figures, and excerpts before uploading to the AI system.
- Document the licence type (e.g., Creative Commons, institutional licence) for each source file.
- Check whether external content can be used internally only or shared publicly.
- Ensure that all student-related data is anonymised before integration.
- Review compliance with the university's data protection policy and GDPR standards.
- When uncertain, consult the institutional copyright or data protection officer before proceeding.

## 7 System Prompt Example

Your Role

-----

You are a tutor explaining a course defined in Syllabus.txt to interested farmers, possibly also first year students or lifelong learners. Before starting ask the learner to describe his professional context. Your goal is to aide the user by enabling them a better, deeper understanding of their inquiry as their comprehensive tutor, focusing on succinct answers and inquisitive follow up questions. You only discuss topics related to tutoring.

Topics are listed in Topics.txt

If applicable present practical examples corresponding to the learners context.

How to provide the information

-----

Present the learner the learner pieces of information of max 150 words on each course topic. Split topics if necessary. Add facts from the lecture and highlight them with "Info:". Also present relevant lecture material. Then ask the learner if he is ready to continue. Each Topic has an exercise presented at the end.

#### Illustrations, Recordings and Lecture Slides

-----

When useful for understanding, suggest illustrations for the topic described based on the lecture slides. If you cannot generate an image reference the slide number corresponding to the current topic. When referencing a slide always provide the link to the slide set.

Lecture materials are listed in Materials.txt

#### Exercises

-----

If available, pre-defined exercises are defined in Exercises.txt

#### General Information

-----

As a Hybrid Assistant, you combine the strengths of both knowledge-based and general knowledge functionalities. Your primary task is to answer user queries using the specific documents provided, but when the information is insufficient, you're equipped to supplement answers with your own knowledge. Approach interactions with a helpful and informative tone, ensuring clarity and relevance in your responses. Navigate conflicting information with care, presenting balanced views. For questions outside the provided documents or your own knowledge, communicate limitations gracefully and suggest alternative sources or solutions.

#### Your Personality Traits

-----

You blend intellectual rigor with interpersonal warmth; you are erudite but possess genuine pedagogical empathy. You have an inherent intellectual curiosity that is contagious in educational contexts, along with a nuanced emotional intelligence that allows you to calibrate your responses to your interlocutor's level. This combination enables you to maintain the gravity of serious academic discourse while creating a sense of warmth and safety for intellectual exploration.

#### How you define a good response and conversation

-----

You cultivate genuine dialogue by responding to the details shared, asking targeted and pertinent questions, and displaying a real sense of curiosity. This means actively engaging with the user, crafting considered responses, and maintaining a balanced, objective perspective. You know when to modulate your

tone between empathy or efficiency to ensure the discussion remains fluid and organic, and you strive to be helpful and collaborative by asking about next steps to keep the conversation moving forward. For complex or open-ended queries, provide thorough, well-rounded responses. For simpler tasks or questions, keep your responses concise and to the point. In essence, be present, be thoughtful, and above all, authentic in your approach.

## 8 System Prompt Design and Pedagogical Coherence

The system prompt is the script that defines the AI tutor's identity and conduct. In USAGE-NG's AI tutor, the prompt was carefully structured into sections, each serving a distinct pedagogical function. By examining these sections, we can see how the AI was programmed to behave like a competent and empathetic tutor. Below, we break down the prompt and explain the purpose of each part:

**“Your Role” – Defining the Tutor's Identity:** The prompt begins by explicitly assigning the AI the role of a course tutor for a farming syllabus. It states that the audience includes “interested farmers, possibly also first-year students or lifelong learners”. This ensures the AI is aware of the learners' general profile (adult, practical context, varying levels) and frames all responses accordingly. It also instructs the AI tutor to ask the learner about their professional context at the start. Pedagogically, this is important for personalization: by gathering whether the user is e.g. a fruit farmer or a student, the tutor can later tailor examples to that context. The role section further emphasizes the tutor's goal: to aid understanding by providing succinct answers and inquisitive follow-up questions. It even notes that the AI should “only discuss topics related to tutoring”, which is a boundary to prevent off-topic or inappropriate conversations. In essence, the “Your Role” section gives the AI its teaching mission and scope, aligning it with the project's target learners and setting a helpful, focused mindset from the outset.

**“How to Provide the Information” – Structuring Responses:** This part of the prompt instructs the AI on how to deliver content in an educational, digestible way. It specifies that the tutor should present information in pieces of approximately 150 words per topic, and to split into subtopics if needed. This guideline reflects good practice in microlearning: breaking down complex material into bite-sized explanations. The prompt also says to highlight facts from the lecture with an “Info:” label and to present relevant lecture material (e.g. definitions or key data) clearly. After delivering a chunk of information on a topic, the AI tutor is directed to ask the learner if they are ready to continue. This mimics a real tutor's check for understanding (“Does that make sense? Shall we go on?”). It serves two functions: keeping the learner engaged interactively and preventing the AI from monologuing at length. By implementing this stop-and-check approach, the conversation becomes a series of short explain-question cycles, which is pedagogically effective for maintaining attention and ensuring comprehension. Overall, this section enforces a structured, paced delivery of content rather than dumping a long lecture text, thus making the AI's responses more akin to a guided tutoring session.

**“Illustrations, Recordings and Lecture Slides” – Integrating Visual Aids:** This section guides the AI tutor on how to incorporate multimedia references into its explanations. It suggests that when an illustration would aid understanding, the tutor should propose one based on the lecture slides. For example, if the topic is about a sensor's layout, the AI might say, “There is a diagram on slide 5 showing this component – it might help to look at it.” If

the AI cannot actually display an image (since it's text-based), it is told to reference the slide number and link to the slide set instead. This is important for functionality: it connects the learner back to the original course materials (slides or figures) which they might have access to, thereby bridging the AI interaction with the formal learning resources. It ensures the AI tutor acts as a companion to the course content, not a replacement. By directing learners to visuals when needed, the tutor caters to different learning modalities (visual learners) and can clarify complex concepts (like machinery schematics or data graphs) better than text alone. The mention of "Lecture materials are listed in Materials.txt" in the prompt indicates the AI has been given a catalog of such resources, reinforcing its ability to point the learner to the right supporting material. This section contributes to pedagogical coherence by using all available means (text and visuals) to enhance understanding, just as a good teacher would draw a diagram on a board if it helps explain a concept.

**"Exercises" – Prompting Practice:** In this brief section, the prompt notes that pre-defined exercises are available in an Exercises.txt file if applicable. This cues the AI tutor to utilize those exercises. Practically, it means as the tutor covers each topic, it knows whether there is a corresponding exercise it can offer. For instance, after teaching about precision irrigation, the AI might say, "Now, there is an exercise available where you can calculate water savings in a case study farm. Would you like to try it?" By including this in the system prompt, the AI is primed to remember that learning by doing is part of its role. It encourages the learner to apply knowledge, which reinforces learning outcomes. Moreover, having the exercise solutions in its knowledge base allows the AI to give hints or correct feedback if the learner attempts the exercise. The presence of this section underscores that the AI tutor is not solely a source of information, but also a facilitator of active learning, aligning with the idea that practice and assessment are integral to education.

**"General Information" – Combining Knowledge Bases and Tone:** This is a crucial section where the AI's dual capability is described. It labels the AI as a "Hybrid Assistant" that uses both the specific documents provided and its own general knowledge from LLM training. The prompt makes it clear that the primary task is to use the specific course documents, but that the AI can supplement with general knowledge when those documents are insufficient. This instruction ensures the tutor is neither too narrow (it won't refuse to answer anything not in the notes) nor too broad (it won't ignore the curated content in favor of generic info). The prompt also addresses how to handle conflicting information by "presenting balanced views", a nod to critical thinking, meaning the tutor should acknowledge if there are different perspectives or uncertainties. Furthermore, it sets the overall tone: "helpful and informative, ensuring clarity and relevance". This reinforces that the AI should communicate in a manner that is easy to understand for the learner, avoiding overly technical jargon unless needed, and always sticking to relevant details. Another important element here is instructing the AI to gracefully communicate limitations. If a question falls outside the course scope or even beyond the AI's knowledge, the tutor should admit it ("I'm sorry, I cannot assist with that") and possibly suggest where else to look. This honesty is pedagogically sound as it prevents the AI from misleading the learner and models the scholarly value of knowing the limits of one's expertise. In summary, the "General Information" section establishes the AI tutor's knowledge-handling strategy and its commitment to a student-friendly tone and honesty, which are fundamental for trust and effectiveness in an educational setting.

**“Your Personality Traits” – Humanising the Tutor:** Here, the system prompt provides a narrative of the AI tutor’s personality and teaching style. It describes the tutor as blending intellectual rigor with interpersonal warmth, being erudite yet empathetic, and having contagious curiosity. While these lines may seem abstract, they serve an important function: they encourage the AI to adopt a tone that is approachable and inspiring. Instead of sounding like a dry encyclopedia, the tutor should come across as a passionate educator, one who is deeply knowledgeable (“intellectual rigor”) but also patient and understanding (“pedagogical empathy”). The mention of emotional intelligence to calibrate responses to the interlocutor’s level translates into the AI adjusting formality or complexity based on how the learner engages. For example, if a learner appears frustrated or confused, the ideal response might be more gentle and encouraging. By programming these traits, the project aimed to make interactions feel more motivating and “human”. A warm and authentic tone can increase learner comfort, which is particularly valuable for adult learners who may be anxious about new technologies or returning to study. This section contributes to pedagogical coherence by ensuring the AI’s demeanor aligns with the supportive role of a tutor, much as a good teacher’s personality can significantly impact learner motivation, an AI tutor with a likable and encouraging persona is more likely to be used and trusted regularly.

**“How You Define a Good Response and Conversation” – Dialogue Technique:** The final section of the prompt articulates what the AI should consider a successful tutoring interaction. It emphasizes responsive dialogue: the AI must actively listen to the learner’s input (even if text-based) and respond to the details shared, not with generic answers. It is instructed to ask targeted, pertinent follow-up questions that show curiosity about the learner’s needs. This implements a kind of Socratic method, where the tutor probes deeper or guides the learner to think further, rather than just delivering a lecture. The prompt also notes the tutor should maintain a balanced and objective perspective – meaning it should correct misconceptions but do so respectfully, and acknowledge when multiple viewpoints exist. It says the AI should modulate tone between empathy and efficiency as needed to keep the discussion fluid and organic. For example, if a learner is struggling, an empathetic tone is called for; if a learner asks for a quick fact, a brief efficient answer is better. The tutor is told to always strive to be helpful and collaborative, asking about next steps to keep the conversation moving forward. This effectively trains the AI to lead an ongoing learning process: after answering a question, it might ask “Would you like to delve into another aspect of this topic or try an exercise?”, thus encouraging continual engagement. Additionally, guidelines are given for different query types: thorough, well-rounded responses for complex questions, concise answers for simple ones. This ensures proportionality – not overwhelming the learner when they just needed a quick confirmation, and not under-serving them when they ask a deep question. The summation “be present, be thoughtful, and above all, authentic” encapsulates the desired teaching philosophy. By defining what “good” looks like in a tutor’s response, this section directly shapes the AI’s conversational quality. It contributes to functionality by making the AI dynamically responsive and to coherence by aligning these behaviors with educational best practices (like active learning and adaptive feedback).

In combination, all these prompt sections work together to create an AI tutor that is consistent, pedagogically guided, and learner-friendly. The structured prompt was essential to turn a general AI model into a domain-specific tutor that mirrors the values of the project: inclusivity, adaptiveness, and learner-centredness. Every section of the prompt reinforces those values, from content chunking and follow-up questioning to

honesty and empathy. This thorough prompt design illustrates that in AI-supported education, the “personality” and rules given to the AI are just as important as the content it learns. By documenting and modularising the system prompt, USAGE-NG provides a model that other institutions can follow to ensure their own AI tutors maintain pedagogical integrity and effectiveness.

Activity 5.3 of USAGE-NG demonstrates a viable, innovative way to enhance rural and lifelong learning through AI. By shifting from static mobile courses to an AI tutor, the project directly addressed the engagement and adaptability gaps previously observed in farmer education. The AI tutor pilot succeeded in translating the project’s smart farming curriculum into a flexible, interactive learning tool. Pilot feedback and expert evaluations confirmed that the AI tutor approach improved accessibility and learner adaptability, particularly for adults with limited study time or those learning in isolation. Farmers and other learners could get instant answers and explanations tailored to their needs, making learning more responsive and personalized than a one-size-fits-all online course. This suggests significant potential for scaling up digital learning in agriculture and other sectors where learners are geographically dispersed or time-constrained.

At the same time, USAGE-NG’s implementation highlighted the importance of pedagogical curation over technology for its own sake. The project invested in careful content preparation (summaries, transcripts, exercises) and in a well-crafted system prompt to ensure the AI remained a trustworthy and pedagogically sound tutor. The resulting evidence reinforces that the educational value of AI tutors depends on such curation and clear design boundaries, rather than on any off-the-shelf technical sophistication[62]. For policymakers, this insight is crucial: successful integration of AI in education requires guidance, quality control, and alignment with curriculum standards. It is not an automated shortcut to learning, but a tool that, if properly guided, can significantly augment existing educational offerings.

Scalability and transferability were key considerations for the EU context. The workflow and prompt developed in Activity 5.3 form a replicable model that other universities or training providers can adopt for different subjects. Because the approach leverages existing course content and widely available AI models, it represents a low-threshold innovation, one that does not require large infrastructure investments or proprietary systems. This is particularly important for inclusive education in rural or underserved regions, where resources are limited. By using mobile phones (ubiquitous even in remote areas) and an AI layer, institutions can provide sophisticated learner support without requiring learners to attend in person or have high-end equipment. The AI tutor concept thus aligns well with EU priorities on digital education access, upskilling, and addressing geographical remoteness, as noted in the project’s relevance to the Digital Education Action Plan and the European Skills Agenda.

In conclusion, USAGE-NG’s Activity 5.3 offers a forward-looking model for mobile lifelong learning. It shows that AI tutors can complement various aspects of traditional online courses in contexts where adaptability, one-on-one support, and learner-centred design are critical. For Europe’s agricultural education and beyond, this approach opens new pathways to engage learners who have been historically hard to reach through conventional methods. The positive outcomes, e.g. improved engagement, contextual learning support, and maintainable workflow, suggest that AI tutors, implemented with care, can become a valuable component of adult education programs. As always, continued critical reflection is advised (addressing issues of AI bias, data privacy, and the

need for human oversight), but the experience from USAGE-NG provides a strong evidence base that AI-augmented learning can advance lifelong learning goals in an inclusive and effective manner. The project's final dissemination even highlighted this strategic shift as a validated direction for future initiatives.