PAIHE

Practice of Artificial Intelligence in Higher Education
Vol. 1 (2026) · Article 2

PRACTICE REPORT

The Practice of Reflective Teaching with AI

Future Skills Development with Offline LLMs

Harald J. Bolsinger[1], Sven Warnke[2], Tanja Mühling[2]

[1] Technical University of Applied Sciences Würzburg-Schweinfurt, THWS Business School, Germany

[2] Technical University of Applied Sciences Würzburg-Schweinfurt, Faculty of Applied Social Sciences, Germany

AT A GLANCE

Context	Approach	Outcome
Final-year bachelor's course on values-oriented corporate communication (media management, THWS), summer term 2026; six student groups	Group process designed to start without AI; several LLMs per group, including offline models; AI use documented, reflected on and assessed	Work products of unusual depth and quality from the instructor's perspective; AI competence enhancement; next run: adapt content

Abstract

This practice report describes how generative AI was integrated into a course on values-oriented corporate communication at the end of a bachelor's program in media management in the summer term of 2026. It documents a first teaching cycle under a minimal consensus for designing university teaching with generative AI, which is set out in the foundational contribution to PAIHE; the course followed an earlier version of that consensus. Students worked in groups on responsibility analyses and values-oriented communication strategies. The group process was designed so that students first noted their own assessments and then used several large language models, including models running offline on their own devices, to add new perspectives and roles; students documented and evaluated their AI use in their presentations, and this documentation entered the assessment. The report explains the didactic rationale, the selection of tools that non-specialists can install and run locally, and the AI rules of the course. From the instructor's perspective, the work products showed unusual depth and quality. These are observations of work products, not evidence of competence; whether the design builds disciplinary and AI-related judgment competence is a question for further teaching cycles. Practical AI competence and understanding was enhanced by continuous use and reflection of usefulness. The report closes with lessons for the next cycle, including the need to shorten and adapt content.

Keywords: generative AI, reflective teaching, future skills, offline large language models, AI literacy, constructive alignment

Published: 5 November 2026
Procedure: Invited contribution by the editors in the launch procedure of the first issue; jointly conceptualized and endorsed by the editors; no review by uninvolved persons.
AI use in the review: AI-generated reviews (Claude, Anthropic) were used as a list of possible objections; all decisions were made by the editors.

Suggested citation: Bolsinger, H. J., Warnke, S., & Mühling, T. (2026). The practice of reflective teaching with AI: Future skills development with offline LLMs. *Practice of Artificial Intelligence in Higher Education, 1*(1), Article 2. https://doi.org/10.83397/PAIHE.2026.02

DOI: https://doi.org/10.83397/PAIHE.2026.02 · **License:** © 2026 Harald J. Bolsinger, Sven Warnke, Tanja Mühling. This article is licensed under CC BY 4.0 (https://creativecommons.org/licenses/by/4.0/).

1 Reflective Teaching with AI

1.1 Didactical Foundations

This article describes one course from a first teaching cycle under the minimal consensus for designing university teaching with generative AI that the foundational contribution to PAIHE sets out (Warnke et al., 2026); as noted there, the course followed an earlier version of the consensus (Warnke et al., 2026, Section 3.3). The consensus aims at a target competence that combines disciplinary competence – here, analyzing responsibility and designing values-oriented communication strategies – with AI-related judgment competence: deciding on AI use, checking AI outputs, integrating them into one's own position, and reflecting on AI use and usefulness. Group work and the comparison of several large language models (LLMs), both used in the course, are extension options of the consensus, not core requirements.

Constructive alignment (Biggs, 1996; Biggs et al., 2022) is the planning framework of this consensus: teaching is planned from the intended learning outcomes, and learning activities and assessment are aligned with them (see also Hailikari et al., 2022). Whether and how AI is used in a task is part of this planning, not a requirement for every course. In the course described here, AI served as an aid to reflection, not only as a tool for producing or enhancing content. Its expected value lies in iterability, availability during the work process, and irritation potential: AI can make implicit assumptions and blind spots visible, but it can also reinforce them unnoticed, which is why its output must always be examined critically. Feedback generated by AI complements feedback from teachers but does not replace it (Weidlich et al., 2025).

For this reason, students worked with several large language models. Comparing the outputs of several AI tools is an extension option of the minimal consensus (Warnke et al., 2026, Section 4.3): it can deepen students' checking of outputs, but it presupposes access to several tools, and its effect has not yet been studied. Algorithmic bias is well documented for AI systems in education (Baker & Hawn, 2022); whether models from different providers or regions systematically reproduce different perspectives is, however, an open empirical question. We therefore treat such differences as a question to be examined, not as a premise. Comparing outputs is intended to train the critical evaluation of AI outputs, a core element of AI literacy (Long & Magerko, 2020), including attention to possible bias.

Generative AI performs well on tasks at the lower levels of the revised Bloom taxonomy, such as remembering and understanding, and is considered capable at higher levels too (Lubbe et al., 2025). For teaching, the decisive question is therefore not what AI can do but which thinking students must do themselves. Where disciplinary knowledge is itself a learning outcome, students have to acquire it, because they need it to check AI outputs and to form their own judgment (Warnke et al., 2026, Sections 2.2 and 2.3). The group process in the course was therefore designed so that students first noted their own assessments, and the analysis itself remained their task. The AI literacy framework of the Digital Education Council (2025) offers a broader, complementary perspective on the competences addressed.

Students document the type and extent of their AI use, including tools, prompts, and an assessment of the benefits, in their presentations and reflections. This serves academic integrity – transparent handling of AI use is increasingly common in the guidelines of leading universities (Moorhouse et al., 2023) – and it creates a record of practical experience for later work such as a thesis. Together, these elements are meant to turn AI use from a cognitive

shortcut into a deliberate, documented practice. Whether they do so is a design conjecture in the sense of the minimal consensus, which further teaching cycles are meant to examine.

1.2 Starting with Practical Experience and AI Diversity for Skill Building

Because students often know only a few AI applications, the first step is to show the diversity of generative AI from providers in different countries. This is useful in every course setting to get the discussion about AI going and to make students curious about each other's AI use. After collecting already known online applications from the student community and discussing individual use scenarios, this list can be supplemented with additional, yet unnamed online applications (see, for example, the supplementary overview of models) to illustrate how varied the global offering of generative AI already is. Here it becomes clear that AI competence is far more than simply mastering digital technology; it also raises cultural, values-oriented, ethical, and geopolitical questions. During the discussion of individual application experiences, it frequently becomes apparent that students, despite the copyright restrictions, upload teaching material to AI services that use it for internal training, because most of them lack paid accounts that allow them to opt out of training. At this point it is advisable to discuss the traceability and storage issues of such uploads in depth. Either some students then propose the use of an offline AI as a solution, or the instructor brings this alternative into the discussion. The variety of available large language models can then be expanded by introducing freely available models. Ideally, instructors report on a few models they themselves use offline, whether on a mobile device, a desktop computer, or a laptop. It is important to also address the investment required to create freely available LLMs to illustrate the revolution and historically unique opportunity in AI availability for personal as well as scientific and commercial use. LLMs that have been developed with substantial research and development funding are freely available for everyone to use, which bridges to the criteria of a "technology for everyone" and its limits.

In this context, it is commonly observed that students occasionally ask how the implementation looks in practice, so that they can set up a comparable configuration on their own devices. If that does not happen, teachers can ask whether anyone is interested in an offline AI setup for learning and research that does not pass course materials or personal data to third parties. Running AI locally does not, however, settle questions of copyright, data protection, or licensing; these have to be checked for the selected model, the data, and the use case.

2 A Toolset that Works for Everyone

2.1 The Evolution of the "Offline AI Use for Teaching" Idea

To autodidactically train the proper use of AI technology without compromising the competencies that have been necessary so far, one of the authors privately purchased a computer in August 2023 that would in the future enable local offline AI applications for testing in teaching. At that time, the device that was perfectly suited for that purpose cost around €3,500, featuring an Intel i9-13900KF CPU, 64 GB DDR5 5600 MHz RAM, an NVIDIA GeForce RTX™ 4090 24 GB GDDR6X, a 1200 W power supply, and elaborate water-cooling. Today, the same configuration costs considerably more. While in previous years computer prices had generally fallen and more performance could be obtained for less money, this relationship has now turned around. Even the development of demand for

powerful AI components and memory chips, all of which rely on the rapid spread of AI applications and the resulting need to equip data centers for the operation of generative AI (inference), already shows this. Since 2022, the AI accelerator chip H100 from Nvidia, which is a specialized General Purpose Computation on Graphics Processing Unit (GPGPU), has become the central bottleneck product for generative AI (Bauduin, 2023; Hiltscher, 2023). Nvidia's GPUs have been continuously increasing in price ever since. Especially at the beginning of the AI revolution, many universities also relied on generative AI from frontier models, which are operated in the external data centers of large providers. In 2024, the Technical University of Applied Sciences Würzburg-Schweinfurt (THWS) introduced campus-wide use of ChatGPT (Technische Hochschule Würzburg-Schweinfurt, 2024). Through the HAWKI platform as an interface to ChatGPT from OpenAI, data-protection compliant use was also sought for teaching. Numerous students and colleagues used this offer and the direct offers from other frontier-model providers such as Anthropic, Google, Meta, Mistral, Perplexity, xAI, and others, more or less in a manner that complied with copyright protection, and especially on the student side very often without considering what data was actually entrusted to the interfaces of externally hosted AI service providers. In conversations with students, it quickly became clear that only very few IT experts truly brought an understanding of how the technology worked and how the input data were actually handled. Added to this is the difference between paid versions with professional features compared to the widely available free access. The period between 2022 and 2025 was, at least in the environment of people in teaching who are not IT-centric, marked by a frightening AI-application dichotomy. On one side were countless users who, as before, simply clicked through all terms and conditions without reading them and then did what was possible. On the other side, a broad group of people condemned the use of AI in the teaching environment because old, knowledge-transfer-based didactic concepts seemed no longer to work. Even experts in training on the topic of generative AI in teaching omitted many open questions about software licenses, data protection, and secure teaching use, and limited themselves to demonstrating what could be done with API connections and the popular AI tools of frontier-model providers.

For this reason, in 2023 we began searching for a combination of tools that could be meaningfully used in our own teaching, which can be deployed offline on already existing everyday devices, operated and understood by non-technicians, and are legally unobstructed for direct use by anyone, and optionally also commercially scalable. We call these tools "Technology for Everyone". Whether a tool is legally unobstructed has to be checked separately at each level: the interface software, the model and its weights, the data entered, and the hosting. Especially in the entrepreneurship area, tools are required for teaching that can also be immediately used commercially, so that the development of AI solutions does not remain a dry exercise. The first choice, given the legal framework conditions and the requirements for research with sensitive data at Bavarian universities, is the targeted offline deployment of LLMs that can access local data and are fully operated on the personal devices that students bring to class. Also concerning the confidentiality of trade secrets, operating generative AI on one's own premises is the first choice. Accordingly, experimenting with own installations was the goal of the hour. The search for a combination of tools was consciously not outsourced to IT specialists. The idea was simple: If a digitally inclined business ethicist such as Harald Bolsinger can get a suitable system up and running quickly and safely without external support, it is ripe enough to be used by other non-IT professionals in teaching, too.

On the aforementioned personal computer, this search began with the testing of Ollama (see https://ollama.com), which came to market in 2023, and numerous other LLM tools for desktop use, up to Nvidia's own tool ChatRTX (at its release still called "Chat with RTX"), which was available in 2024 only for very powerful Nvidia GPUs (Bastian, 2024). At the beginning of the search it took a full two days until the prerequisites for running an offline LLM were created. Under Windows 11, a Windows Subsystem for Linux 2 had to be installed, then RTX Windows drivers updated, then the CUDA toolkit installed in Windows Subsystem for Linux, then Ollama and then the further tested tools with Nvidia NIM set up, which are not named here because of their tested abundance. In the meantime, this configuration process has become fully automated and the tools considered are now mature enough for teaching use. Thus, what were experimental applications became a "Technology for Everyone" suitable for use in teaching beyond IT experts.

2.2 Pragmatic Software Selection: Technology for Everyone Criteria

As of 6 October 2026, the authors have tested 60 different LLMs both offline and online and provide excerpts in a supplementary overview of models to inspire further testing by the readers of this publication on their own. From the current perspective and based on experience from summer term 2026, especially very small, quantized offline LLMs – also sometimes called small language models (SLMs) – are suitable for teaching use, as they also run on weak and older devices. Recommendations and background information for teaching based on a specific teaching use-case are given in the following section. The LLMs alone are not enough – you also need a platform providing an interface for them that makes inference happen. A software package suitable for teaching use has to meet conditions that the authors would describe as "technology for everyone". We consider the following features necessary for long-term meaningful use in teaching:

– Simple installation, including easy download of curated LLMs
– Possibility of self-hosting on personal devices
– Hardware performance requirements that are met by common current computers and even older ones
– Multilingual variants available for all major operating systems and architectures
– Released under a free license (e.g. MIT license), commercial use also possible
– Integration of freely usable offline LLMs as well as cloud LLMs (up to frontier models) via API
– Integration of source material like PDFs with local data storage (without internet access) with an easy-to-use RAG implementation
– Ongoing maintenance and further development

Additional nice-to-have features for use in special learning environments are:

– Implementation of agent system functionality for automating recurring tasks
– Optional internet access and Model Context Protocol implementation
– Options for online self-hosting
– Options for paid online hosting in the cloud of the tool provider

The last two options allow a small, isolated start on a laptop or even a mobile phone and let students later scale to a self-hosted server (e.g., with Docker images) or a ready-to-use cloud environment with many users or potentially the whole world. This is relevant for making

students develop their own solutions that can be rolled out into the real world, as needed in practice-oriented courses about entrepreneurship.

Ultimately, a very well-suited basic tool for operating offline LLMs on one's own computers that meets the above criteria is AnythingLLM from the US company Mintplex Labs Inc., founded in 2022 by Timothy Carambat. He earned a bachelor's degree in mechanical engineering at Southeastern Louisiana University in 2016 and learned intensively autodidactically to develop software like AnythingLLM. AnythingLLM meets all the above criteria, is available for Windows ARM and x64, macOS Silicon and Intel, Linux, and even as a mobile app for Android with connectivity/synchronization options to the other versions. With its own API, AnythingLLM can also be made available in the classroom directly to students whose hardware is too weak, so that they can use other people's hardware. Since version 1.16, it usually takes only three steps to get it running, thanks to the automated installation process: download and install AnythingLLM, an LLM suitable for the specific hardware is automatically chosen in the process and downloaded, and you then can start chatting. It now runs on nearly all devices without difficulty, and later versions (1.18 as of October 2026) have brought further improvements and features.

Within the course, sufficient time must be reserved so that all students install this basic tool on their device and successfully receive a response to an initial prompt. Students can help each other with individual installation problems under the guidance of the instructor and the additional use of online generative AI for special installation questions. After successful installation, nothing technically stands in the way of the use of offline LLMs in future learning activities.

Without goal-oriented rules this learning setting cannot succeed, which is why the instructor must establish and explain these rules in advance for the use of any AI around the course. In the best case, there are general rules of the university or at least of the faculty that address AI use in courses; where there are none, you should start making your own set of rules. As an example, Appendix A contains the rule set of one of the authors as inspiration for creating own rule sets.

2.3 Recommendations for LLM-Knowledge and LLMs for Teaching

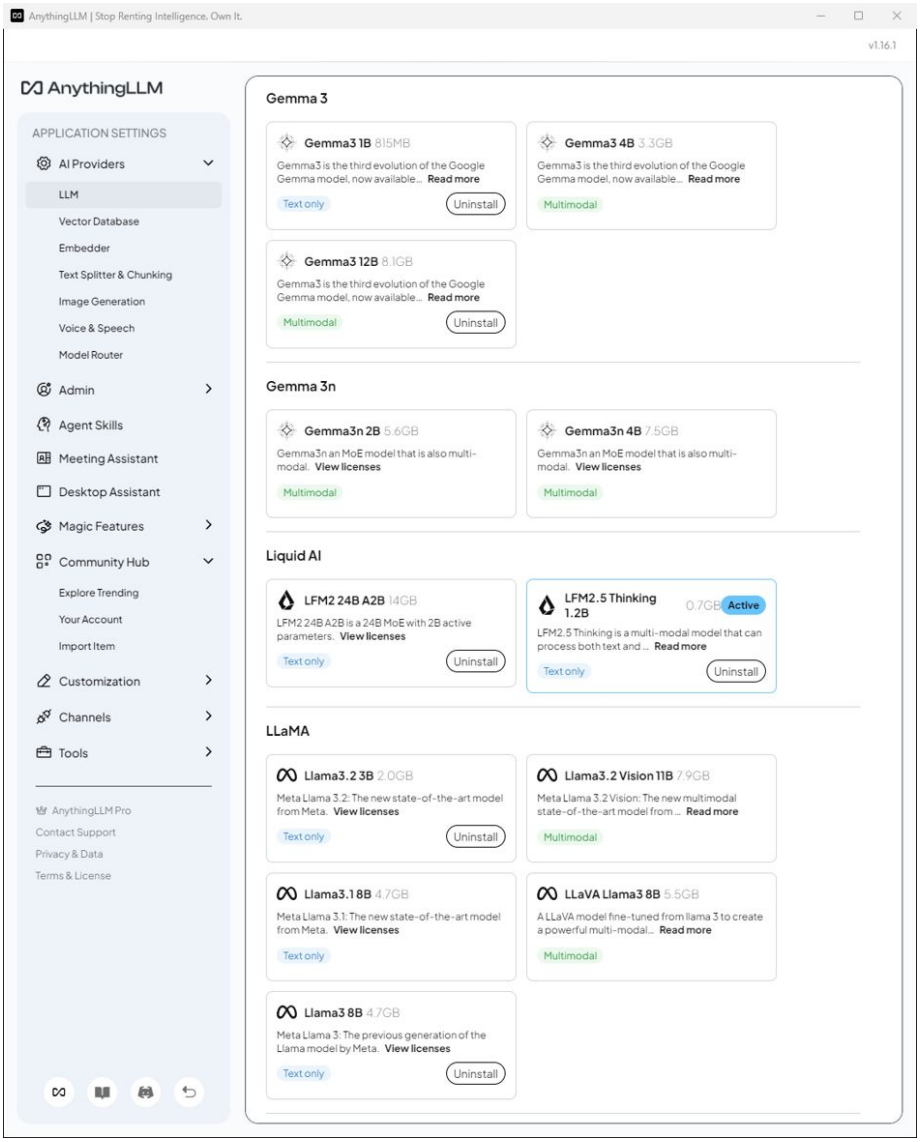
LLM basic knowledge is essential for all disciplines, not only to understand the technical limits and pitfalls of applying LLMs in one's own field but also to competently assess the geopolitical, cultural, and ethical backgrounds. LLM basic knowledge is already excellently produced as learning material in numerous variants. Students and instructors alike must keep themselves up to date autonomously. A very good, free introduction with 21 units is offered by Microsoft at https://github.com/microsoft/generative-ai-for-beginners. It is highly recommended because it is also available in several languages at https://github.com/microsoft/generative-ai-for-beginners/tree/main/translations and can be easily linked into one's own teaching material. It can be expected that Microsoft will keep the content up to date, as they ultimately aim to facilitate with the learning material the use of LLMs in their own digital ecosystem. For use in one's own teaching, the focus can be on chapters 01–05 and 15, 16. These cover terminology that is important for understanding and selecting LLMs, such as Generative Pre-trained Transformer, Open-Weight/Open-Source, Service versus Model. Responsible use of AI is also addressed in an easily understandable way, always with application-oriented explanations, e.g., how Retrieval Augmented

Generation (RAG) and Vector Database can be integrated, up to using the online platform Hugging Face (https://huggingface.co), which NVIDIA acquired recently, as an LLM provider.

When selecting an LLM it is advisable to use as small as possible Q3 or Q4 quantized models with AnythingLLM that can potentially run on as many different devices in a course as possible. This is the case when the model fits entirely in the graphics card memory for Windows PCs or in the RAM of Apple PCs with Apple-Silicon-Chips or also in the RAM of mobile devices that are suitable for on-device AI-inference. Small models also run acceptably if there is no dedicated graphics card but sufficient RAM or when there is a Neural Processing Unit (NPU) in the device. To get a reliable assessment of which models will work well on which hardware, it is recommended to use websites like https://canitrun.dev or ask any online LLM and then download the working GPT-Generated Unified Format (GGUF) files from Hugging Face for import into AnythingLLM. GGUF is the standard format for local AI inference and pretty easy to import. The simplest step is, however, to use the integrated download function of AnythingLLM for pre-selected LLMs. In the example of AnythingLLM version 1.16.1 (Figure 1), options include Google Gemma3 1B at 815 MB, Liquid AI LFM2.5 1.2B at 0.7 GB, or Llama3.2 3B at 2.0 GB (or not shown in the image: Qwen3 0.6B at 0.6 GB) that all also run on reduced hardware configurations.

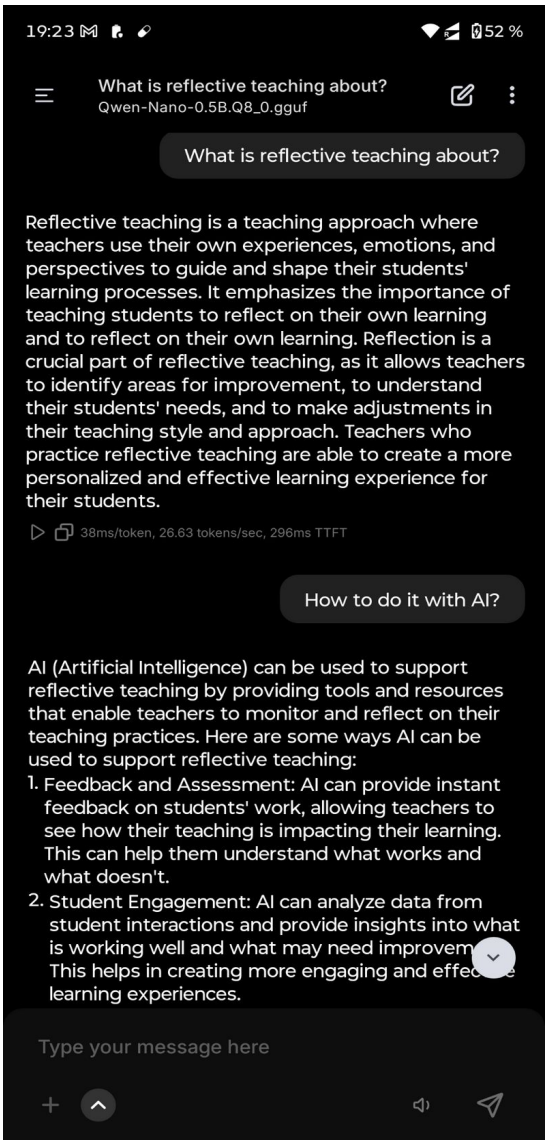
Figure 1

Pre-Selected LLMs for Hassle-Free Download in AnythingLLM (Version 1.16.1)



For use on iPhones as well as Android mobile devices with sufficient RAM of more than 8 GB, the free app PocketPal AI at https://pocketpal.dev is also suitable. AnythingLLM mobile is also available and has integrated the download of LLMs from Hugging Face and many more features since version 1.3, but is currently available only for Android devices (iOS support is planned). When using mobile devices, too, it makes sense to have one application for all mobile platforms in the classroom. So until AnythingLLM mobile also runs on iPhones, PocketPal AI is, however, an option for the more advanced users with good devices and can also import nearly any GGUF file from Hugging Face (Figure 2). There is also a leaderboard at https://pocketpal.dev/leaderboard that gives you a reliable assessment of which LLMs will work well on which hardware. Especially small models such as those named above run very smoothly on many devices.

Figure 2
PocketPal AI on Android with Qwen-Nano-0.5B Q8 (525.12 MB)



The training data used, and thus the worldview of the curators of training data and the creators as well as the implemented safety features, weights, and system prompts, influence the responses of every LLM. Added to this is the problem that hallucinations caused by LLMs cannot be technically eliminated completely, but can only be limited. Even though system prompts can protect against problematic outputs, it is impossible to remove political influence from the respective LLMs. Freedom of opinion, press, and religion, minority rights, and human rights in general are often systematically restricted by political motives. In China, for example, there are numerous regulations that, among other things, aim to ensure that communist values are not undermined by Chinese LLMs. For advanced instructors and relevant research, comparisons of outputs from LLMs whose safety features have been removed or heavily restricted are also an interesting use-case. These are identifiable on Hugging Face with the labels ‘heretic’, ‘uncensored’, or ‘abliterated’. However, these are not recommended for testing live with students in courses, as the outputs can be very problematic and should not be made visible without prior assessment by the responsible instructors. Some examples are listed in the supplementary overview of models for completeness, because students can and will also independently access and use these

LLMs. It is important to at least be aware of these options and then be able to address the corresponding issues.

3 A Case Study at THWS Business School

3.1 Course Topic and Scientific Framing

The course “Values-oriented Corporate Communication in the Social Market Economy” (a flagship course at the end of the Bachelor in Media Management at THWS) in the Summer Term 2026 served as a resonant space for the first full implementation of reflective teaching with offline AI. Harald Bolsinger received the best teacher award 2026 from Technical University of Applied Sciences Würzburg-Schweinfurt (THWS) for that course, in the digital teaching category, and it also stands out because of its exemplary character within the THWS Business School. The course’s overarching goal is to enable students to develop and implement effective communication strategies. Of particular importance is the cultivation of an understanding of and respect for diverse value orientations and worldviews, the development of a personal stance on key values, and the ability to analyze responsibility and trust in communication surroundings professionally and contextually. Students learn to recognize and address ethical dilemmas that arise from multiple rationalities and ultimately to develop goal- and values-oriented communication based on specific norms and values catalogues. This supports restoring lost trust in people and organizations, signaling responsibility, and taking it on. The content-related learning outcomes and the assessment format (portfolio assessment) were not changed by integrating reflective teaching with AI; only the path to competence development was broadened. In addition, the AI rules of the course state AI-related learning outcomes: students should use AI reflectively as a complement to their own thinking, check AI outputs for accuracy, relevance, and bias, and document their AI use transparently (Appendix A). These outcomes entered the assessment through the documentation of AI use, through an assessment criterion that explicitly addressed the practical, comparative use of several AI tools and the identification of their strengths and weaknesses for the task at hand, from the perspective of media managers, and through the discussion during the presentations, in which students had to show the depth of their understanding, both of the subject and of their AI use. Oral colloquia were held with all students. Using AI was required for parts of the assessed work; the disciplinary learning outcomes could nevertheless also be achieved without AI, and free options were available (Appendix A). From the instructor’s perspective, the use of AI increased the depth and quality with which the learning outcomes were achieved.

Shifting the course required a change in perspective for the instructor. Instead of lecturing most of the time, the instructor consistently adopted the role of a coaching and learning companion, discovering, further developing, and updating the teaching content together with the students. Each subject unit had to be re-thought with meaningful questions that, when answered, contribute to the above professional competence goals. Instead of classical lecturing and knowledge explanation, the students are faced with questions that they have to answer in groups together using at least the available teaching material and AI in a way that can be presented in plenary sessions. These questions were designed so that, through active group work, students can reach appropriate results within a suitably chosen timeframe.

3.2 AI Integration

The basic introductory knowledge of how to use AI tools was collected in the first lecture through a poll in the class. The vast majority used and were only familiar with the most common AI tools from U.S. providers. These were used like a common search engine, often without any understanding of legal frameworks, prompting strategies, or the resource costs in terms of energy and water consumption. Students rarely used AI under guidance within the context of courses; they used it only on their own initiative and as a starting point for generating results. In many other courses, AI use was forbidden, and where it was not, students reported that there was no in-depth practical guidance on what is useful and how to use it. Of course, exceptions prove the rule – occasionally there were voluntary special lectures on AI that provided an in-depth practical introduction to the world of generative AI options for research and scientific writing and graph building. Almost no one in the affected course had attended them, though. Offline use on personal devices with freely available LLMs was still done by no one.

To be able to work with AI in a meaningful and legally secure way, all students were required to independently complete the e-learning course “Introduction to AI at THWS” with valuable basic information about AI functions that are allowed at THWS and how you can correctly use artificial intelligence in your daily work or studies at THWS. This e-learning course helps students to use AI more competently, responsibly, and confidently. Students had to upload the competence certificate earned there into the course room of the corresponding lecture to get bonus points. In addition, AI rules for the lecture (Appendix A) were drafted and discussed in detail with the students at the beginning. Already in the second lecture, AnythingLLM was explained and run on students’ devices. After that, groups were formed so that each group had sufficiently powerful hardware available. For special cases where no device was available, an exception was made to use https://academiccloud.de/services/chatai/, where selected LLMs are hosted for academic use. The entire course content was structured in blocks as follows (Table 1):

Table 1
Course Schedule, Summer Term 2026

Date	Time	Topics	Format
18 March 2026	10:00–17:00	General topics, **AI use in comparison**, ethics in corporate communication	In person
8 April 2026	10:00–17:00	Group formation, **AI didactics and technical prerequisites**, worldview, values, morality and conflicting ethical types, **AI in corporate communication I (reflective teaching with AI – first steps)**	In person
6 May 2026	10:00–17:00	**AI in corporate communication II (reflective teaching with AI)**, fundamentals of responsibility	In person
13 May 2026	10:00–17:00	**Analyzing and demonstrating responsibility (reflective teaching with AI)**	In person / hybrid online
27 May 2026	13:30–17:00	Trust and communication – arguing effectively and appropriately for the target group	In person / hybrid online
10 June 2026	10:00–17:00	Sustainability: materiality analysis, responsibility analysis (students' presentations)	In person
17 June 2026	10:00–17:00	Human rights I (case study Gao-Feng), **AI-enhanced group work on responsibility analysis and trust building with coaching**	In person
24 June 2026	13:30–17:00	Guidelines, audits, human rights II, **social market economy as a framework for freedom in entrepreneurial activity (reflective teaching with AI)**	In person / hybrid online
8 July 2026	10:00–17:00	**Group presentations of responsibility analyses and AI usefulness**, teaching discussions, survey on reflective teaching with AI	In person

Note. Topics with an AI focus are set in bold.

The students were divided into six groups. Each group received its own Conceptboard (https://conceptboard.com) for collaboration (also with students who participated in the respective course hybrid or asynchronously) and for documenting the results. Each group was asked to use up to three different AI tools, including at least one offline LLM. Online use of LLMs was only allowed in accordance with the AI rules of THWS and the course. For each group, the assignment was to use LLMs from at least two different "cultural backgrounds" – in practice, from providers based in different countries – and, at the very end of the process, to search for and verify sources on subject-specific questions with Scopus AI (https://www.elsevier.com/en-gb/products/scopus/scopus-ai). Scopus AI is itself a generative search tool: its results help to identify potentially relevant sources, which then have to be checked against the original publications.

– Group 1 was asked to use Google Gemini, Mistral Le Chat, and Apertus.
– Group 2 was asked to use ChatGPT, DeepSeek, and Teuken.
– Group 3 was asked to use Microsoft Copilot, Kimi, and Qwen.
– Group 4 was asked to use Claude, Qwen, and Llama.

– Group 5 was asked to use Grok, LFM, and Gemma.
– Group 6 was asked to use Perplexity (the Sonar standard model from Perplexity itself), Meta AI, and GLM.

LFM2.5 is developed by Liquid AI, which is based in the USA; all tools assigned to group 5 therefore came from US providers. Assigning the tools to the groups ensured that a range of tools was compared, but it limited students' own justified choice of tools, which the minimal consensus counts as part of AI-related judgment (Warnke et al., 2026, Section 2.3).

To avoid cognitive offloading and address as many future competencies simultaneously as possible, the group sessions followed a standardized process that had been intensively discussed and refined by the authors beforehand (Figure 3).

Figure 3
The Process of Reflective Teaching with AI



Note. Created with the help of ChatGPT 5.2 (OpenAI).

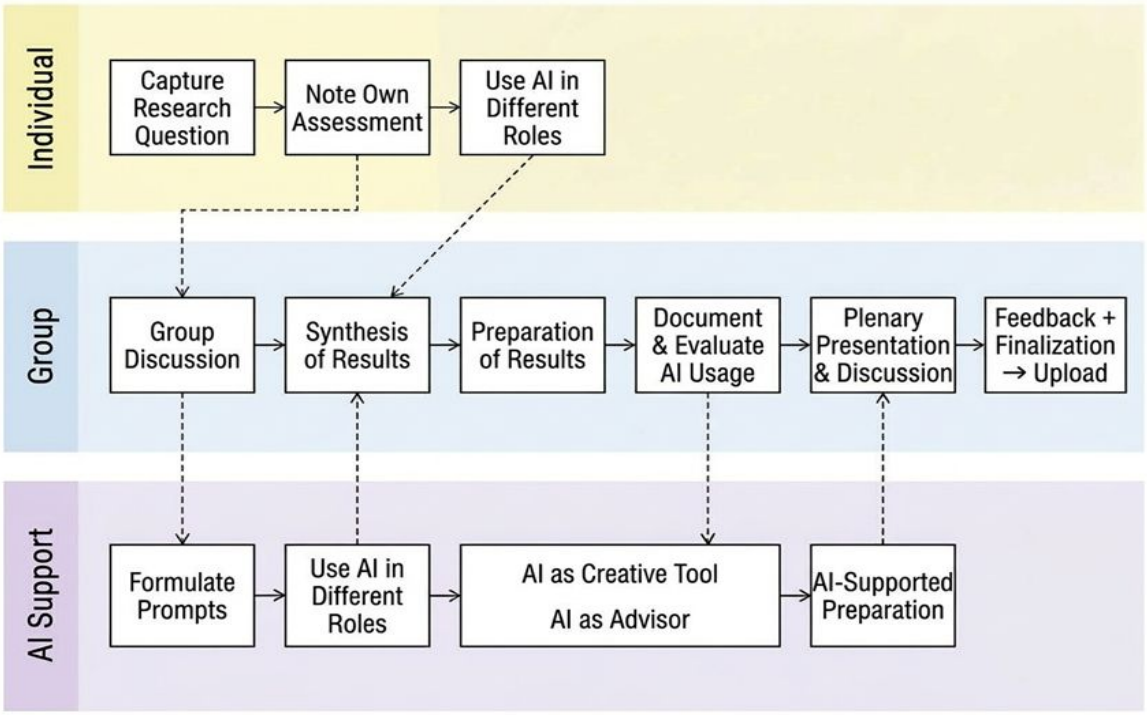
The use of AI was therefore woven into a group process that was designed to start without AI and that addresses both critical thinking and communication skills. Students were free to

organize their work; some already used AI while the questions were being explained, as they were used to doing, and occasionally brought these results into the group. In the group process, missing perspectives and roles for answering a question were identified and only then specifically addressed with AI. The core was thus the group process, which was complemented and improved in content by AI, with AI functioning as a creativity tool, a reflective aid, and a tool for preparing the results in an attractive format that can be presented (Figures 3 and 4).

Measured against the minimal consensus (Warnke et al., 2026, Section 4.2), the first core design function, “own prior work before using AI,” was built into the process but not implemented consistently. The second, “AI as feedback on one’s own interim work,” occurred occasionally: groups developed different strategies for using AI in the process, and some asked the AI to comment on their own assessment and decided which comments to take up. The third, “reflecting on, justifying, and documenting AI use,” was implemented: groups recorded their reflection on AI use on their Conceptboard and in their presentations; much of the reflection took place orally in the discussion during the presentations, and these discussions influenced how they used AI afterwards.

Figure 4

Individual, Group, and AI Tasks in the Learning Process



Note. The figure shows which tasks were performed individually, in the group, and with AI support. Individual AI use before the group discussion was permitted. Created with the help of gamma.app.

Reflective teaching sessions with AI were typically characterized by a brief thematic introduction by the instructor, followed by an explanation of the questions to be addressed. All groups worked on the same questions, which meant that a variety of perspectives on the same topic emerged in the plenary presentations.

Below we provide an example of the questions from a one-day session that aimed to develop a clear and scientific understanding of responsibility analyses:

- How can responsibility be conceptually defined in corporate communication?
  - How do legal, moral, social, and communicative responsibilities differ?
  - Which actors are involved, and how does the use of AI influence the responsibility chain in corporate communication?
- What should an analytical instrument look like that allows responsibility in case studies to be assessed fully and comparably, in order to derive communicative measures?
  - What questions must always be clarified? How are relevant stakeholders identified?
  - How can responsibility relationships be made visible and fully taken into account?
  - How are all responsibility categories considered in analyses?
- What criteria allow an ex-post analysis of whether responsibility was actually taken on or merely communicated symbolically?
  - How can responsibility for consequences be systematically reconstructed?
  - What role does the foreseeability of possible effects play?
  - What role do stakeholder expectations play?

3.3 Core Experiences and Outlook

Especially questions like “What should an analytical instrument look like?” changed teaching enormously, as the instructor did not just deliver his instruments for use by the students; the students worked out and reflected on their own instruments before applying them. That led to many new aspects and automation options with AI, for example in stakeholder analysis, that one would not have dared to dream about a few years ago.

In the third meeting, students already presented first results from their group work and put them up for discussion. The quality of the elaborations in terms of content and design varied widely. The instructor clarified that the use of AI made top-class quality possible and from then on accepted only work that met this standard, as professional practice also demands. Presentations also had to show which parts were the group’s own analysis, how AI outputs had been checked, and which AI tools had been used for what. Students carried out each step of the responsibility analysis themselves – developing the analytical instrument, identifying actors and stakeholders, assigning responsibility relationships to the corresponding norms, and justifying their assessment – and used AI in every step for support, including research and the preparation of presentations. In the process, they became aware of the limits of AI use and remained in the driver’s seat. From the instructor’s perspective, AI helped them to make progress and to go into considerably more depth than would have been possible without it; the group presentations on a comprehensive responsibility analysis of a corporate scandal and on a communication strategy for building trust showed a depth and quality not seen in earlier runs of this course. These are observations of work products. They cannot by themselves show what students have learned (Warnke et al., 2026, Section 1) without adding topical discussion with the students about specific points of the presentations.

Plenary presentations of the results, including experiences with the AI tools and the self-chosen approach to usage, followed only after all questions had been fully addressed within a predetermined time frame. The variety of the presentations was another plus point: it led to solutions that went beyond the content pre-selected by the instructor. The instructor also gained valuable impulses for further improvement and updating of the teaching material, which will be taken into account in the next iteration in terms of the type of questions.

Although there were occasionally incorrect or uninformed points in the presentations, these were directly addressed by the instructor's quality assurance, so that the professional limits of the AI tools used were also made clear. Not least for this reason, scientific quality assurance by the instructor remains strictly necessary for the presentations.

To describe how students report their AI use at the beginning and at the end of AI-enhanced courses, and to adjust the teaching formats accordingly, the authors jointly developed a questionnaire. Its structure and first findings are reported in a separate article in this issue (Mühling et al., 2026). Reliable statements about the effects of this form of teaching require accompanying research with suitable designs.

In terms of content, the course must be shortened because the form of teaching involves significantly more group processes, presentations, and discussions, and the introduction to AI tools and the enabling of concrete AI use also require additional time. Some of this can be compensated with flipped classroom approaches, but the overall workload of the respective course offering must be kept in mind.

Overall, the course produced work of a quality and range that the instructor had not seen in earlier runs, and students presented their experiences with the AI tools in the plenary sessions. These observations show that the design is feasible in this setting. Whether it supports the intended disciplinary and AI-related judgment competence cannot be inferred from the work products; it is a question for further teaching cycles. AI will continue to be used in the instructor's future courses.

Statement on AI Use

We used generative AI in preparing this article.

Literature search. ChatGPT 5.2 Pro (Deep Research) and Perplexity (Education Pro): license information in the supplementary overview of models; Gemini, ChatGPT, Perplexity, and Copilot: information on training data.

Drafting and revising. Claude (Anthropic), October 2026: drafts of the abstract and of revised passages in Sections 1.1 and 3.1–3.3 and in Appendix A. We checked and revised all wording.

Translation. GPT-OSS 20b MXFP4 (OpenAI) in AnythingLLM: translation into English of large parts of the text, which were written in German, using the system prompt "You are a business ethics scientist and translator. Translate all the given text to English. Do not alter names and citations. Return only your translation."

Figures and tables. ChatGPT 5.2 (OpenAI): Figure 3; gamma.app: Figure 4. We checked both.

Other tasks. AI Policy Creator for Teaching Staff, version 2 (University of Bamberg): first version of the AI rules in Appendix A, which we then revised manually.

We have checked all content and take responsibility for it.

Funding and Conflicts of Interest

No external funding was received for this article. The authors are the editors of PAIHE. This constitutes a conflict of interest, because in the launch procedure of the first issue no uninvolved persons reviewed the article. No other conflicts of interest exist.

Author contributions (CRediT): Harald Bolsinger: conceptualization, investigation, resources, visualization, writing – original draft, writing – review & editing. Sven Warnke: conceptualization. Tanja Mühling: conceptualization.

Statement on Research Ethics

This practice report is based on the instructor's observations and on course materials. The poll in the first session was anonymous; students had been informed beforehand that it was part of the authors' research project on AI use in teaching. No data on individual students are reported. The questionnaire administered at the end of the course is reported, including its ethics statement, in a separate article in this issue (Mühling et al., 2026).

References

Baker, R. S., & Hawn, A. (2022). Algorithmic bias in education. *International Journal of Artificial Intelligence in Education, 32*(4), 1052–1092. https://doi.org/10.1007/s40593-021-00285-9

Bastian, M. (2024, February 2). Nvidias „Chat with RTX“ ist ein kostenloser, lokaler Daten-Chatbot [Nvidia's "Chat with RTX" is a free, local data chatbot]. *heise online*. https://www.heise.de/news/Nvidias-Chat-with-RTX-ist-ein-kostenloser-lokaler-Daten-Chatbot-9627934.html

Bauduin, S. (2023, August 16). KI statt Gaming: Nvidia soll 2023 mehr als 500.000 HPC-Beschleuniger absetzen [AI instead of gaming: Nvidia expected to sell more than 500,000 HPC accelerators in 2023]. *PC Games Hardware*. https://www.pcgameshardware.de/Kuenstliche-Intelligenz-Hardware-279517/News/Nvidia-verkauft-eine-halbe-Million-KI-Beschleuniger-in-2023-1426781/

Biggs, J. (1996). Enhancing teaching through constructive alignment. *Higher Education, 32*(3), 347–364. https://doi.org/10.1007/BF00138871

Biggs, J., Tang, C., & Kennedy, G. (2022). *Teaching for quality learning at university* (5th ed.). Open University Press.

Digital Education Council. (2025, March 3). *Digital Education Council AI literacy framework*. https://www.digitaleducationcouncil.com/post/digital-education-council-ai-literacy-framework

Hailikari, T., Virtanen, V., Vesalainen, M., & Postareff, L. (2022). Student perspectives on how different elements of constructive alignment support active learning. *Active Learning in Higher Education, 23*(3), 217–231. https://doi.org/10.1177/1469787421989160

Herrmann, D. (2026). *KI-Policy Presets* [AI policy presets]. https://ki-policy.org

Hiltscher, J. (2023, November 28). 500.000 KI-GPUs in einem Quartal sind noch zu wenig [500,000 AI GPUs in one quarter are still not enough]. *Golem*. https://www.golem.de/news/h100-a100-500-000-ki-gpus-in-einem-quartal-sind-noch-zu-wenig-2311-179820.html

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). ACM. https://doi.org/10.1145/3313831.3376727

Lubbe, A., Marais, E., & Kruger, D. (2025). Cultivating independent thinkers: The triad of artificial intelligence, Bloom's taxonomy and critical thinking in assessment pedagogy. *Education and Information Technologies, 30*(12), 17589–17622. https://doi.org/10.1007/s10639-025-13476-x

Moorhouse, B. L., Yeo, M. A., & Wan, Y. (2023). Generative AI tools and assessment: Guidelines of the world's top-ranking universities. *Computers and Education Open, 5*, Article 100151. https://doi.org/10.1016/j.caeo.2023.100151

Mühling, T., Warnke, S., & Bolsinger, H. J. (2026). Reflektierte KI-Nutzung in Lehrveranstaltungen: Befunde aus dem Pretest am Ende des Sommersemesters 2026 [Reflective use of AI in university courses: Findings from a pretest at the end of the 2026 summer semester]. *Practice of Artificial Intelligence in Higher Education, 1*(1), Article 3. https://doi.org/10.83397/PAIHE.2026.03

Technische Hochschule Würzburg-Schweinfurt. (2024, July 1). *Technische Hochschule Würzburg-Schweinfurt führt hochschulweit ChatGPT ein* [Technical University of Applied Sciences Würzburg-Schweinfurt introduces ChatGPT university-wide]. https://www.thws.de/service/news-presse/pressemeldungen/thema/technische-hochschule-wuerzburg-schweinfurt-fuehrt-hochschulweit-chatgpt-ein/

Warnke, S., Bolsinger, H. J., & Mühling, T. (2026). Leistung ist nicht Lernen: Ein Minimalkonsens für die Gestaltung von Hochschullehre mit generativer KI [Performance is not learning: A minimal consensus for designing university teaching with generative AI]. *Practice of Artificial Intelligence in Higher Education, 1*(1), Article 1. https://doi.org/10.83397/PAIHE.2026.01

Weidlich, J., Gotsch, F., Schudel, K., Marusic-Würscher, C., Mazzarella, J., Bolten, H., Bütler, D., Luger, S., Wohlfender, B., & Maag Merki, K. (2025). Teacher, peer, or AI? Comparing effects of feedback sources in higher education. *Computers and Education Open, 9*, Article 100300. https://doi.org/10.1016/j.caeo.2025.100300

Appendix A: AI Rules for Teaching (Current Version)

The rules applied in the course were essentially the same. The passage on access to AI tools was revised after the course; notes marked “In the course described in this article” refer to this course only.

These rules set out the guidelines for the use of AI in the described course and all courses by one of the authors, Harald Bolsinger. Its purpose is to create clarity for students and to ensure fair assessment, while at the same time allowing students to explore the possibilities of AI. This framework was first generated with the help of AI Policy Creator for Teaching Staff Version 2 (operated by the Chief Information Office of the University of Bamberg) in March 2026 and manually revised on 16 September 2026. Generative AI was also used as an auxiliary tool in the drafting and translation of the texts. The newest version of the AI Policy Creator is available at https://ki-policy.org and supports English and German (Herrmann, 2026).

1 Fundamentals and Learning Objectives

Basic principle regarding the use of AI

The use of AI tools (such as ChatGPT, Copilot, Claude, Gemini, HAWKI, etc.) is encouraged in this course as a means of supporting the learning process. However, the specific provisions of this policy regarding permitted and prohibited use, disclosure, and academic integrity apply.

Explanation: AI is part of everyday life and work. In this course, a transparent and reflective approach to AI is to be encouraged as a matter of course. This general permission does not exempt you from the obligation to disclose its use and to comply with the principles of good academic practice. Your own contribution and your critical thinking always remain the priority and are also the competencies that must be developed in order to assess AI-generated results.

Learning objectives with regard to AI

In this course, you will deepen the following fundamental AI-related competencies:

– Reflective use of AI tools as a complement to your own thinking and as support for learning processes.
– Effective formulation of prompts to achieve precise and helpful results.
– Critical evaluation of AI-generated content with regard to accuracy, relevance, and possible bias.
– Transparent and correct documentation of AI use in an academic context in accordance with good academic practice.

In addition to these basic AI competencies, in this course we will jointly explore:

– The specific strengths and weaknesses of different AI models for different tasks.
– Strategies for reviewing and validating AI-generated information using scientific methods.
– The ethical implications of using AI in the discipline.

Explanation: Working with AI requires specific skills. These learning objectives are intended to help you not only use AI, but to do so critically, effectively, and responsibly in an academic context.

2 Permitted and Prohibited Use

Permitted use without disclosure requirement

The following types of AI use are permitted as personal learning aids and do not need to be specifically disclosed:

– Use for individual exam preparation (e.g. creating practice questions, summaries of your own notes, flashcards).
– Generating explanations of concepts or technical terms for your own understanding.
– Translating foreign-language texts (e.g. journal articles) for your own understanding.
– Use in exercise phases during the course that are explicitly intended for this purpose (when guided by the instructor).
– Spelling and grammar checking.

Explanation: If you use AI solely for your own learning (e.g. to understand texts or to prepare for exams), without this directly feeding into an assessed submission, you do not need to indicate it.

Permitted use with disclosure requirement

The following types of AI use are permitted, provided they are disclosed transparently and traceably in accordance with the requirements of this policy:

– Support in brainstorming, structuring, and outlining written assignments, homework, or presentations.
– Help with the linguistic revision and improvement of your own texts (beyond pure spelling/grammar correction, e.g. style, wording suggestions).
– Generating examples, analogies, or alternative explanatory approaches to illustrate your own arguments.
– Brainstorming possible arguments, counterarguments, or perspectives for an analysis.
– Support with literature research.

The decisive point is: The core achievement (analysis, argumentation, critical classification, your own conclusions) must be produced by you yourself.

Explanation: If AI supports you in tasks that are assessed (e.g. brainstorming, outlining, text improvement), you must disclose this. The purpose is to make transparent where the support came from. However, the actual thinking and responsibility for the content always remain with you.

Guidance on the disclosure requirement follows in Section 4 of this policy.

Prohibited AI use

The following types of AI use are expressly not permitted in this course:

– Submitting AI-generated content as your own work without adequate disclosure — this may be regarded as an attempt to deceive.
– Using AI to circumvent learning objectives that explicitly require independent critical thinking, creative problem-solving, analytical skills, or specific methodological competencies.
– Uploading copyrighted course materials or personal data to cloud-based AI systems without the explicit written permission of the rights holders (see Section 3, “Copyright and Data Protection”).

Explanation: AI may not be used to pretend that work is your own or to bypass learning objectives. Submitting AI-generated texts as your own work without disclosure is an attempt to deceive and will generally result in failing the assessment.

3 Copyright and Data Protection

Handling copyrighted materials

The teaching materials provided in this course (slides, scripts, assignments, etc.) are protected by copyright. You are not permitted to upload these materials, or substantial parts of them, to external cloud-based AI services (such as ChatGPT, Copilot, Claude, Gemini, HAWKI, etc.) or to enter them there for analysis or processing.

Explanation: Course materials, like textbooks, are the intellectual property of instructors or the university. At later stages of development of AI tools, it will generally become traceable where which content originated. Uploading course materials to external services may constitute unauthorized distribution and may violate copyright law. Therefore, follow the instructions carefully regarding which materials, if any, you may use with AI tools.

Data protection aspects

When using AI services, pay attention to data protection:

– Do not enter any personal data (e.g. names, matriculation numbers, email addresses, discussion posts) into external AI tools — neither your own nor those of other students or teaching staff.
– Be aware that many cloud-based AI services store your inputs and may potentially use them to train their models.
– Where possible, use more privacy-friendly alternatives such as the university’s own AI service or AI models running locally on your computer.

Explanation: Protecting your data and the data of others is important. External AI companies may not be subject to the strict European data protection rules. Therefore, avoid disclosing personal information. Your inputs may be stored and reused.

4 Disclosure of AI Use

Format of disclosure

Please disclose the use of AI systems in written submissions as follows:

“Documentation of AI Use” in the appendix:

– Tools used: List all AI tools used, including name, provider, and, if applicable, version (e.g. “ChatGPT-5.6 Terra (OpenAI, as of Sept 2026)”).
– Usage log (examples): Document example prompts and the resulting AI responses that influenced your work. Briefly describe how you used or revised the response. (Extent by arrangement with the instructor, e.g. 2–3 examples.)
– Reflection: Add a short personal reflection (approx. 100–200 words): How did the use of AI influence your work process? Where were the advantages, and where were the challenges or limitations?

Direct disclosure in the text (only in cases of direct adoption):

– If you adopt AI-generated passages verbatim or with only slight paraphrasing, mark these directly in the text (e.g. by means of a footnote: “This paragraph is based on a suggestion by ChatGPT-5.6.”).

In addition, the general declaration of independent work may also need to be submitted in accordance with the regulations of your examination rules.

Explanation: Transparency is crucial for academic integrity. The disclosure requirement ensures that your own contribution remains comprehensible and that it is clear which tools you used. The format specified here is intended to document the use of AI fairly and consistently.

Exceptions to the disclosure requirement

Explicit disclosure of AI use is not required in the following cases:

– For individual preparation for examinations (e.g. creating practice questions, summaries).
– For use solely to promote understanding (e.g. explanations, translations for personal use).
– For ungraded exercises, unless otherwise specified by the instructor.
– For the use of standard spelling and grammar checks.

Explanation: Not every use of AI must be documented. If you use AI purely as a personal learning aid, without this directly feeding into an assessed piece of work, the disclosure requirement does not apply.

5 Equal Opportunity and Access

Access to AI tools

In order to ensure equal opportunity, this course does not assume that you have access to paid premium AI services. No paid AI subscription is required for successful participation. The disciplinary learning objectives can be achieved without AI; for activities that develop and show AI-related judgment, free options are available:

– Use of free basic versions of common AI providers (e.g. ChatGPT, Claude, Gemini).
– The HAWKI service, which is available free of charge to all students.

The use of paid services is at your own responsibility and at your own expense.

The course and the assessment format have been designed in such a way that the potentially higher quality of assessed work resulting from the use of premium services does not have any significant impact on the grading of the assessment.

To ensure this goal reliably, a spontaneous oral subject discussion may be required in addition to a written paper.

Explanation: Not everyone has the same access to expensive AI tools. Therefore, it is ensured that you can complete all tasks using free or university-provided tools and that you will not be disadvantaged in grading if you do not use premium versions.

Assessment Principles and Use of AI

We are aware that the use of AI in unsupervised assessments (e.g. term papers) can hardly be controlled and may endanger equal opportunity and fairness. In order nevertheless to ensure a valid assessment of your individual competencies, the following measures are applied in this course:

– Accompanying iterative result-generation processes within the course.
– Written work and presentations may be supplemented by an oral colloquium in which understanding and independence are reviewed.
– The assignment formats place greater emphasis on transfer tasks, critical analysis, and application, and avoid pure reproduction of knowledge wherever possible.
– Within the chosen portfolio assessment format, process-related elements are included in the grading (e.g. drafts, reflections, active participation).

Explanation: Since it is difficult to control whether and how AI has been used in term papers, etc., we adapt the assessment formats. The aim is to assess your actual abilities and understanding fairly, regardless of whether AI was used in the background. The stated measures (e.g. exam, oral examination, certain types of tasks) help to achieve this.

In the course described in this article, the grade was based on the documentation of AI use, active participation, the presentations and the discussion during them, and bonus points for the e-learning certificate; drafts and reflection texts were not graded.

6 Practical Tips for Using AI

Sustainability

PLEASE CHECK BEFORE USING ANY AI HOSTED IN DATA CENTERS WHETHER IT CAN REALLY ADD VALUE TO CONVENTIONAL WORK METHODS! EVERY AI REQUEST USES COMPUTING POWER AND THEREFORE VALUABLE ENERGY AS WELL AS WATER (FOR COOLING). ALWAYS DEVELOP A CLEAR PLAN FOR HOW AND WHAT YOU WANT TO USE AI FOR BEFORE DIVING INTO IMPLEMENTATION AND PROMPTING! AVOID USELESS FUN REQUESTS, USE OFFLINE AI ON YOUR OWN DEVICE WHEN POSSIBLE, AND REFRAIN FROM CREATING USELESS CONTENT/MEDIA!

Prompt strategies

For better results when using AI, the following strategies may be helpful when formulating prompts:

– Be specific: The more precise your question, the better the answer. Avoid vague prompts.

– Provide context: Briefly explain the background or goal of your request (e.g. “I am writing a term paper on X and need Y”).
– Assign a role: Ask the AI to take on a certain role (e.g. “Explain it to me as if I were a first-year student”, “Answer as a critical reviewer”).
– Specify the format using examples: Tell the AI what the answer should look like (e.g. “List in bullet points”, “Create a table”, “Summarize in three sentences”, “Answer like in this example”).
– Work iteratively: Ask follow-up questions, request clarification, or ask for revision of the answer.
– Phrase neutrally: Avoid leading questions that push the answer in a certain direction (instead of “Is it true that X is good?”, better: “What are the advantages and disadvantages of X?”).

Explanation: The way you “talk” to an AI (prompt it) significantly affects the quality of the answer. With precise and well-considered prompts, you will obtain more useful and relevant results for your studies.

Handling AI errors (“hallucinations,” etc.)

AI systems make mistakes. Do not rely blindly on AI-generated information. Pay particular attention to:

– Factual errors (“hallucinations”): The AI invents facts, dates, or events that sound plausible but are false.
– Invented sources: The AI cites sources (books, articles) that do not exist or misrepresents their content.
– Outdated knowledge: The training data of the AI is not always current. Information may be outdated.
– Incomplete or one-sided presentation: Important aspects, perspectives, or counterarguments may be missing.
– Bias: The AI may reproduce unconscious prejudices from the training data.

Strategies for verification

– Check central claims and facts against academic sources (e.g. scholarly literature, databases) and with science-based AI like, for example, Scopus AI.
– Verify quotations and references — does the source really exist? Is the statement accurate?
– Compare the answers of different AI tools or actively search for opposing viewpoints.
– Be especially skeptical about very specific or surprising claims.

Explanation: AI tools are impressive, but not infallible. They can provide false information or omit important details. Your task as scholars is to critically review this information and not simply adopt it. Responsibility for the correctness of your work always lies with you, not with the AI.

Support for Questions About the Use of AI

If you have questions or uncertainties about the use of AI in this course or about the interpretation of this policy, you are welcome to contact me:

– By email or in the course forum

– Within the course itself (we plan time for discussion)
– During my office hours

Explanation: If you are unsure how you may use AI or how you should disclose its use, please ask. Open communication helps avoid misunderstandings. It is better to ask once too often than to violate the rules unintentionally.

7 Use of AI by Teaching Staff

General remarks

The instructors teaching this course may also use AI tools to create materials or to support communication. In doing so, we always ensure compliance with data protection and copyright law and critically review AI-generated results.

Explanation: We use AI tools to improve and make our teaching more efficient. However, this is always done in compliance with the legal framework and with the aim of providing you with the best possible learning experience.

Course preparation and materials

Where appropriate, we may use AI tools to support the following tasks:

– Planning the structure and content of the course.
– Creating and revising teaching materials (e.g. slides, summaries, visualizations).
– Generating practice tasks or case studies.
– Optimizing assignments or descriptions.

All content created or revised in this way is carefully reviewed by us and selected according to didactic criteria.

Explanation: AI can help make course materials more varied or generate new ideas for exercises. This serves to improve the quality of teaching for you.

Feedback and grading

When grading assessed work (e.g. exams, term papers, presentations) that contributes to your final grade, we do not use AI tools. Grading is carried out exclusively by human examiners.

We do use AI tools to provide you with formative feedback (e.g. on drafts, practice tasks) that is intended to help you learn, but that does not count toward the final grade. The grading of final, assessment-relevant work is carried out exclusively by human examiners. We carefully review the results of AI tools and ensure that no personal data and no copyrighted content are transmitted to cloud-based AI systems.

In the course described in this article, the instructor did not use AI to give feedback; feedback on students' work came from people.