

Policy on students' AI use at the Faculty of Technology

Faculty of Technology, Bielefeld University¹

October 2026

The purpose of this document is to define a joint understanding how teachers at the Faculty of Technology, Bielefeld University, intend to deal with students' use of AI systems. The target audience are teachers; however, in designing the document, the comments of students as well as administrators have been considered and continued feedback is invited and welcome. The document is not to be understood as binding law (the binding regulations for teaching are still defined by laws, examination regulations, subject-specific regulations, module descriptions, etc.). Rather, it formulates a goal and standards that are to be translated into regulations and teaching practice.

The overarching goal of this policy is to foster students' autonomy (Münzer, Go & Paaßen, 2026) and intellectual sovereignty (Wissenschaftsrat, 2026) in a deep sense: not only in terms of letting students choose whether and how they wish to employ AI tools but, more importantly, regulating AI use such that students acquire the prerequisite knowledge and skills to use AI tools in a competent and responsible fashion, remain fully in control of the interaction instead of ceding autonomy or offloading cognition. As scientific/factual basis of this policy, we recognize that:

1. Many students wish to use AI tools to help them while some students remain opposed (Hoffmann and Schmidt, 2026).
2. Automatic detection software for AI-generated text (or code) is unreliable such that other methods have to be applied (Gehring and Paaßen, 2025).
3. The usage practices for AI tools are highly diverse, ranging from full generation of solutions over revisions to feedback and targeted questions (Hüsch, Horstmann and Breiter, 2025).
4. Some usage practices harm student learning and lead to de-skilling (Lehmann, Cornelius and Sting, 2025).
5. To decide when and how to use AI tools and to competently check their output, knowledge and skills in the subject – i.e. computer science, biotechnology, math, or related fields in our faculty - as well as meta-knowledge and skills (AI literacy) are required (Denny et al., 2023; Perry et al., 2023).
6. In the spirit of the European AI act as well as good scientific practice it is mandatory to be transparent that and how AI tools have been used (DFG, 2023).

Overall, we recognize that different didactic contexts (introductory versus advanced courses, theoretical versus applied courses, exercises versus projects versus essays versus exams, different subjects) require different regulations to best foster student autonomy. At the same time, we aim for standardization to make it clearer for students which rules apply and easier for teachers to decide their course of action.

¹Main authors of first draft: Benjamin Paaßen, Ellen Baake, Thomas Hermann

As a compromise between flexibility and standardization, we arrive at a **three-tier “traffic light system” regarding AI use: Green permits AI tool use broadly, yellow for specific purposes, and red forbids AI use.**

In general (and in alignment with other rules and regulations) we define AI systems as tools, not as co-authors or sources.² That is, we do not distinguish between unpermitted AI use or other ways to cheat (such as hiring ghostwriters or copying solutions from other students). Consequently, it is important to point out that, even at the green tier, **general rules** always apply (just as road traffic regulations still apply even if a traffic light is green), with particular emphasis on three principles: *transparency*, meaning that the use of AI as a tool (“Hilfsmittel”) has to be communicated and hence made tracable ; *responsibility*, meaning that students need to remain the responsible authors of the work in a meaningful sense; *competency* meaning that test performance (“Prüfungsleistungen”) need to remain proof of having achieved the learning goals.

In the following section, we define the three tiers and provide extensive recommendations how to implement them, namely how to communicate the policy to students (before the course/thesis starts), how to implement the policy in courses, and how violations of the policy may be sanctioned.

Beyond this policy, the Faculty of Technology actively pursues the teaching of AI-related knowledge and skills as well as the pedagogically meaningful utilization of AI tools for teaching but these topics are outside the scope of this particular document.

Green tier: AI tool use is permitted

The green tier permits AI use for any purpose as long as general rules are still adhered to.

Didactic reasoning

The green tier is recommended if students can be assumed to already possess the prerequisite knowledge and skills to competently judge the outputs of AI tools and correct them if necessary. In that case, permitting AI tool use provides students with larger freedom of choice and enables them to practice responsible AI use as well as scaffold their learning with AI support. The green tier is particularly useful in courses where (creative) AI use is part of the learning goal of the course (e.g. a course on software development in collaboration with agentic tools). Nonetheless, AI systems should be used in ways that avoid de-skilling and in line with the three principles of transparency, responsibility, and competency. Hence, tasks and examination formats still need to be designed such that students achieve the learning goals.

Students correctly declaring AI use

While AI tool use is permitted, students are still required to declare their AI use. We recommend to do so separately for each task/exercise where AI tools were used and to specify (at a high level) how AI tools were used. The following phrases may be helpful:

- During the writing of this report, ChatGPT was used to correct typos, grammar errors, and style

²It is acknowledged that there is plenty of research (especially in our faculty) on AI-human teamwork in which AI systems *fulfill the social role* of a co-worker. However, this does not change the legal or ethical status of AI systems.

(but not the content) of the text.

- In this exercise, Cursor was used to generate code for plotting the results.
- In this exercise, BIKI (with Gemma-4 26b) was used to provide feedback for my solution.

These high-level descriptions have their drawbacks. For example, “provide feedback” could also mean that the LLM provided the correct solution as part of the feedback which was then incorporated in the student’s solution. Nonetheless, in the interest of keeping the declaration requirement light-weight for most students, we recommend a concise documentation style. If the declaration seems ambiguous/imprecise, students can be asked to clarify individually. If teachers notice that the load of individual questions becomes overly large, the declaration requirement can be sharpened to, e.g., include links to full chatlogs.

Special requirements apply when declaring AI use in a thesis. Teachers should clarify the AI regulations (at least the tier) *before* the thesis starts and students must declare AI use in the “Eigenständigkeitserklärung” (Declaration of Own Work) at the end of the thesis to allow for evaluating and grading the student’s performance. A minimal phrasing is:

- In coordination with the thesis reviewers, the following AI tools were using during the creation of this thesis: ChatGPT for stylistic and grammar revisions, Cursor for code debugging, ...

If more detail is needed, the legal department of Bielefeld University provides a template for the declaration³ in which each usage instance is a row and columns are the tool that was used, the purpose of the use, the location of use, and a description of use.

Implementing the policy

While the green tier is most permissive, we still need to check if students comply with the transparency principle (did they declare every use of AI tools correctly?), the responsibility principle (can they explain and defend their work?) and the competency principle (does their work still prove their own competency?).

To enforce the transparency requirement, it may be tempting to rely on automatic detectors whether code or text was AI-written. However, these detectors are not reliable enough for policy enforcement (Gehring and Paaßen, 2025); and the use of detectors would likely lead to a climate of distrust between teachers and students that is incompatible with the spirit of teaching at Bielefeld University and the Faculty of Technology.

Instead, our approach is to check transparency, responsibility, and competency at the same time. In particular, we recommend tests that require students to explain and defend their work without the help of AI tools. We aim to enable at least the following two options:

1. include written tests/quizzes during lectures/seminars/tutorials in which students need to respond to questions regarding their previously submitted exercises and are not allowed to use

³ <https://www.uni-bielefeld.de/themen/pruefungsrecht/chatgpt/Muster-Eigenständigkeitserklärungen-Chatbots.docx> (German)

AI tools;

2. require students to present their solutions to tutors in tutorial sessions.

These measures should be part of the grading scheme, meaning: if students fail to convincingly explain and defend their solutions, points should not be granted. In that way, we clearly make transparent to students that the skill to explain and defend is a crucial part of passing our courses.

Even more, if students fail to explain and defend their solutions convincingly but their submitted solution is correct, this indicates that they have not written the solution themselves. They may have copied the solutions from other students, the internet, or utilized AI tools. Either way, teachers may move toward sanctions.

For Bachelor or Master theses, the Faculty aims to make an oral defense a mandatory and graded part of the examination for theses at the green tier.

Sanctioning violations

In implementing the policy, multiple scenarios can occur in practice that require different reactions:

1. Students declare AI use and can explain and defend their work. In this case, **no action is required**. While it may still be true that students used AI tools beyond regulations, this will be difficult to detect and the risk of false positive detections may be high. However, if the case is obvious (e.g. due to hallucinated content or other clear markers), even when the student can explain the work, it can be sanctioned as a violation of the transparency requirement.
2. Students do not declare any AI use, yet can *not* explain and defend their work. In this case, the students likely did not generate the solution themselves. We don't know whether AI was used or another mechanism – and we purposefully do not distinguish. Either way, this is a case where **sanctions may be in order**.
3. Students do declare AI use and can not explain and defend their work. In this case, **no specific sanctions are required**. Students should just not get any points for the solution.

Whenever case 2 applies, the following mechanism should be employed (always following the binding regulations in §19 of the exam regulations)⁴:

1. Students are informed of the suspicion that they did not author their submitted work and have the opportunity to respond in writing, and optionally in a meeting with a witness.
2. Teachers determine whether they are still sure (i.e., can prove or have strong evidence) that the students did not author their submitted work. If so, the following possible sanctioning choices exist (these do not have to be applied in order; teachers need to decide which option fits best)⁵:

⁴This is adapted from the descriptions of the legal department. More details can be found there (in German): www.uni-bielefeld.de/themen/pruefungsrecht/chatgpt/#comp_000068d206ad_0000003d8a_00af

⁵The actual examination legal rules can be found here (in German): https://www.uni-bielefeld.de/themen/pruefungsrecht/pruefungsrechtliche-rahmen/#comp_00005f7143a4_0000002185_743e

- In very minor cases, such as students forgetting to declare AI use for one task even though they declared it correctly in another task; or students making small, formal mistakes in their declaration: Giving no points for a specific exercise.
 - In cases where students violate regulations for the first time and it can reasonably be estimated that they did not intend to mislead the teacher: giving no points for an entire exercise sheet. In the first semester, a strong warning may also be employed.
 - Enforcing more frequent checks for future submissions or even retro-actively checking prior submissions.
 - In severe cases, such as repeated violations or offloading crucial and substantial parts of the coursework to an AI tool such that the student's own accomplishments are not sufficient to pass the course anymore: let the student fail the entire course and notify the examination office of a "Täuschungsversuch".
3. Students have to be informed in writing (e-mail suffices) of the selected sanction. Some guidance for viable formulations will be provided by the examination office. Communications should be preserved in case needed for later objections.

Explaining the policy to students

The course-specific AI policy must be made very transparent to students early on in a standardized way. This helps to enhance student understanding and standardization across the faculty.

The first step is to include the following phrasing in the ekVV course description: *"In this course, the use of AI tools is permitted provided that you are transparent about the use and can still explain and defend your work. Undeclared AI use may be sanctioned."*

Further, in the first session of the course, we recommend to include a short (ca. 5 minutes) section to explain the AI policy. The following points should be made clear:

1. Announce that this course operates in the green tier. You can re-use the phrasings from the ekVV and moodle. If different tiers apply to different parts of the course (e.g., green to exercises, red to exams) this should also be made clear on the first slide. Public domain traffic light symbols (SVG format) for the slides can be found here: <https://uni-bielefeld.sciebo.de/s/WGfsa33kcWQeKbD>
2. Explain the didactic reasons for choosing this tier (see above for inspiration).
3. Specify how the policy is implemented in the course. In particular, please explain how you check the responsibility criterion (explain & defend), how students should declare AI use, and that violations may be sanctioned.
4. Optionally (but highly recommended): provide a practical example of a task relating to the course and an incorrect/partially correct solution provided by an LLM to discuss with the students what is correct and what is wrong about the solution and to motivate why it is important to acquire domain-specific knowledge and skills to competently judge the output of AI tools and, thus, use AI tools responsibly.

Examples of slides designed to convey these points are: Human-Centric Machine Learning (Seminar + Exercises with written report as examination; green tier): <https://uni-bielefeld.sciebo.de/s/LKE6ZZPMAowZjHZ>

Yellow tier: AI tool use is permitted under restrictions

In the yellow tier, AI tool use is permitted for selected purposes as specified by the teacher (e.g., the purposes of feedback and revision, but not for full generation of solutions). Teachers may also specify that the use of AI tools for certain tasks will reduce the number of points awarded (because the tasks get meaningfully easier and thus demonstrate less skill/knowledge). Otherwise, the same rules as for green apply.

Didactic reasoning

The yellow tier is recommended if teachers want to enable students to use AI tools but students may still lack (some of) the required knowledge and skills to competently judge AI outputs or generate solutions themselves. Further, a didactic reason for the yellow tier may be that AI tools substantially reduce the skill needed to arrive at a correct solution such that lower marks should be granted for a correct solution that involved AI use. The yellow tier provides teachers with a wide range of possible choices how exactly they want to specify and restrict permitted AI use. Therefore, explaining the didactic reasons for the specification is particularly important in this tier to make students understand why the specification is exactly as it is.

Students correctly declaring AI use

We broadly recommend to follow the same declaration rules as in the green tier (see above). However, due to the specific restrictions for AI use applied in the yellow tier, students should be asked to be precise about how they used AI tools to make clear that it was aligned with the policy. As such, the tabular template provided by the legal department of Bielefeld University may be helpful⁶. In this template, each row describes one instance of using an AI tool. The columns contain information which tool was used, the purpose of the use, the location of use, and a description of use.

Implementing the policy

To implement the policy, the same mechanisms are recommended as in the green tier (see above). However, when testing students to explain and defend their work, it can be particularly important to test aspects where no AI use is permitted (e.g. generate code from scratch or generate a mathematical argument).

Sanctioning mechanism

Broadly speaking, the same sanctioning mechanism applies as in the green tier (see above). However, a special scenario is that students do declare AI use that exceeds the restrictions of what is permitted in the course. In this case, sanctions may be considered. It is likely that students misunderstood either the restrictions or the policy (given that they declared the use). If this is a case, lower sanctions may

⁶ <https://www.uni-bielefeld.de/themen/pruefungsrecht/chatgpt/Muster-Eigenständigkeitserklärungen-Chatbots.docx> (German)

be applied (e.g. only striking the points for an exercise). However, repeated declarations of AI use that violates the specifications may be grounds for harsher sanctions.

Explaining the policy to students

The course-specific rules (in line with this policy) should be made very transparent to students early on in a standardized way. This helps to enhance student understanding and standardization across the faculty.

The first step is to include the following phrasing in the ekVV course description (adapt passages in brackets [] as needed): *“In this course, the use of AI tools is permitted [for the following purposes: ...]. [If AI tools are used, less points may be awarded]. You are required to be transparent about AI use and explain and defend your work. Use of AI tools beyond this permission may be sanctioned.”*

If the permission is limited to only parts of the course (e.g. the exercises), this should be specified in place of *“In this course”*. Here is a draft formulation from the course *“Betriebssysteme”* (https://www.ti.uni-bielefeld.de/html/teaching/SS26/techinf2/index.html#ki_tools):

*“Bitte beachten Sie, dass in den Übungen in dieser Veranstaltung der Einsatz von KI-Werkzeugen (z.B. Chat-GPT) für die direkte Erzeugung von Programmcode vollständig **untersagt** ist.*

*Es ist aber bspw. **zulässig**, KI-Werkzeuge zu verwenden, um die Aufgabenstellung besser zu verstehen, um Erklärungen zum Einsatz von Programmierwerkzeugen zu erhalten oder um Grundlagenwissen zu recherchieren, das für die Bearbeitung der Aufgabe notwendig ist. **Unzulässig** ist dagegen die Nutzung von KI-Werkzeugen, um direkt Programmcode (komplett oder in Teilen) generieren zu lassen.*

Die Übungen in der Grundlagenveranstaltung “Betriebssysteme” dienen einerseits der Vermittlung praktischer Fertigkeiten bei der betriebssystemnahen Programmierung, sie sind aber andererseits auch ein unverzichtbares Mittel, um das theoretische Wissen aus der Vorlesung durch den praktischen Einsatz zu vertiefen und zu verankern. Ohne eine aktive Auseinandersetzung mit dem Stoff der Vorlesung würde das Verständnis des Vorlesungsstoffs in großen Teilen oberflächlich bleiben (siehe das Video zu Lernstrategien). Um diese doppelte Rolle der Programmierübungen sicherzustellen, wäre eine Freigabe von KI-Werkzeugen zur direkten Erzeugung der Programme nicht sinnvoll. Selbst wenn der/die Studierende die vom KI-Werkzeug erzeugte Lösung verstehen und erklären könnte, wäre er/sie vermutlich nicht in der Lage, die Lösung selbst zu erzeugen. Damit fehlen wesentliche Aspekte eines erfolgreichen Lernvorgangs wie die eigenständige Überwindung von Schwierigkeiten und die breite Verankerung des Wissens, bspw. durch die Entwicklung und Bewertung alternativer (darunter auch nicht erfolgreicher) Ansätze.”

Here is another example specifying an AI policy-compatible rule for bachelor/master theses:

“Yellow tier: Using AI tools (such as large language models) in the thesis is permitted for the following purposes:

- revising text and giving feedback on text

- as component in the method after discussion with the supervisor, following best practices in the field
- generating code if the functionality is pre-specified by the student and checked
- providing feedback on code, including debugging support

Other purposes are, per default, not permitted but can be discussed with the supervisor (this list will, then, be augmented)."

Further, in the first session of the course, we recommend to include a short (ca. 5 minutes) section to explain the AI policy. The following points should be made:

1. Announce that this course operates in the yellow tier. You can re-use the phrasings from the ekVV and moodle. If different tiers apply to different parts of the course (e.g., yellow to exercises, red to exams) this should also be made clear on the first slide. Public domain traffic light symbols (SVG format) for the slides can be found here: <https://uni-bielefeld.sciebo.de/s/WGfsa33kcWQeKbD>
2. Explain the didactic reasons for choosing this tier (see above for inspiration). This is particularly important for the yellow tier. This slide can also be used to specify more precisely which usage practices are permitted.
3. Specify how the policy is implemented in the course. In particular, please explain how you check the responsibility criterion (explain & defend), how students should declare AI use, how many points are deducted if they do so (if that is the case), and that violations may be sanctioned.
4. Optionally (but highly recommended): provide a practical example of a task relating to the course and an incorrect/partially correct solution provided by an LLM to discuss with the students what is correct and what is wrong about the solution and to motivate why it is important to acquire domain-specific knowledge and skills to competently judge the output of AI tools and, thus, use AI tools responsibly.

We don't have slide examples for this tier yet. However, here is an example website information set up for the course "Betriebssysteme": https://www.ti.uni-bielefeld.de/html/teaching/SS26/techinf2/index.html#ki_tools

Red tier: AI tool use is forbidden

In the red tier, the use of AI tools is forbidden. A main challenge in this tier is to define what counts as AI tool. The boundary will likely shift slightly over time the more experience we (as teachers and students) gather in interacting with AI tools. It is recommended to provide concrete examples to students what kind of technology is considered AI tools so that they know what to avoid (see the "explaining the policy to students" section). The other main challenge is to enforce the policy. In this regard, we recommend essentially the same mechanisms as for the other tiers. More details are provided below.

Didactic reasoning

The red tier is recommended if students cannot be expected to have the required skills to competently judge AI outputs or using AI tools bears severe risk of undermining the learning process. This can specifically apply to courses that build foundational skills for the entire field (e.g. introductory programming) or where the practice of generating a solution yourself is a core part of the skills to be acquired (e.g. in math; scientific writing). While these are strong didactic reasons that are concerned with the autonomy of students in a deeper sense (Münzer, Go, Paaßen, 2026; Wissenschaftsrat, 2026), the red tier still restricts the freedom of choice of students and we, as teachers, should make the rationale for these restrictions clear and that we do not make such decisions lightly (refer also to the “explaining the policy to students” section below). In particular, it may be helpful to announce that, later in their studies, AI tool use will be permitted across many courses; however, we forbid it for the foundational skills.

Students correctly declaring AI use

In this tier, no AI use is permitted such that any declaration of AI use would provide a basis for sanctions.

Implementing the policy

To implement the policy, the same mechanisms are recommended as in the green tier (see above), meaning tests and/or presentations. However, it is still noteworthy that undeclared AI use is hard to detect or to distinguish from other kinds of cheating, such as hiring ghostwriters or copying solutions from other students. Hence, we treat these essentially the same, in line with §19 of the examination regulations.

Sanctioning mechanism

In the red tier, declared AI use or the suspicion that a student did not author their submitted work may both trigger the sanctioning workflow as described in the green tier above. However, special considerations should apply in the first semester of studies. In this situation, students are new to the AI policy at the Faculty of Technology. Hence, first-time violations may be treated with warnings and lighter sanctions (only removing points for one task). However, the warning should be communicated clearly to highlight the importance of the policy and make clear how severe the pedagogical harms are and that undeclared AI use is treated as serious violation of good scientific practice.

Explaining the policy to students

The course-specific AI policy must be made very transparent to students early on in a standardized way. This helps to enhance student understanding and standardization across the faculty.

The first step is to include the following phrasing in the ekVV course description: *“In this course, the use of AI tools is strictly forbidden. Any use of AI tools may be sanctioned.”*

If the restriction is limited to only parts of the course, this should be specified in place of *“In this course”*. Here is an example from “Introduction to Data Mining” where the green tier applies to the exercises and the red tier applies to the exam:

“In the exercises, the use of AI tools is permitted provided that you are transparent about the use and

can still explain and defend your work. Undeclared AI use may be sanctioned. In the exam, the use of AI tools is forbidden. Any AI use may be sanctioned."

Further, in the first session of the course, we recommend to include a short (ca. 5 minutes) section to explain the AI policy. The following points are recommended:

1. Announce that this course operates in the red tier. You can re-use the phrasings from the ekVV and moodle. If different tiers apply to different parts of the course (e.g., yellow to exercises, red to exams) this should also be made clear on the first slide. Public domain traffic light symbols (SVG format) for the slides can be found here: <https://uni-bielefeld.sciebo.de/s/WGfsa33kcWQeKbD>
2. Explain the didactic reasons for choosing this tier (see above for inspiration). This is particularly important for the red tier.
3. Specify how the policy is implemented in the course. In particular, please explain examples of tools that are counted as AI tools (does, e.g., *grammarly* count? Does *DeepL* count? Probably yes, but both should be made clear).
4. Optionally (but highly recommended), provide a practical example of a task relating to the course and an incorrect/partially correct solution provided by an LLM to discuss with the students what is correct and what is wrong about the solution and to motivate why the use of AI tools is forbidden in this course.

A example of such slides for the course "Mathematische Biologie" are linked here: <https://uni-bielefeld.sciebo.de/s/pmCmEaP7CJjwndr> (password: techfak_ai)

Mixed Tiers

In many courses, it will be required to apply different tiers to different parts of the course. Perhaps most common: Exercises operate at the green or yellow tier, but the final examination operates at the red tier. In that case, this should be specified in the ekVV course description and explained to the students. Here is an example from "Introduction to Data Mining" where the green tier applies to the exercises and the red tier applies to the exam.

"In the exercises, the use of AI tools is permitted provided that you are transparent about the use and can still explain and defend your work. Undeclared AI use may be sanctioned. In the exam, the use of AI tools (and any digital tools) is forbidden. Any AI use may be sanctioned."

Here is an example for the slides used in "Introduction to Data Mining" to explain the policy for students: <https://uni-bielefeld.sciebo.de/s/TRf4Rc4xX7PdE5M>

Literature

- Denny, P., Kumar, V., & Giacaman, N. (2023). Conversing with Copilot: Exploring Prompt Engineering for Solving CS1 Problems Using Natural Language. Proceedings of the 54th ACM Technical Symposium on Computer Science Education V. 1, 1136–1142. <https://doi.org/10.1145/3545945.3569823>
- DFG. (2023, September). Stellungnahme des Präsidiums der Deutschen

Forschungsgemeinschaft (DFG) zum Einfluss generativer Modelle für die Text- und Bilderstellung auf die Wissenschaften und das Förderhandeln der DFG.

<https://www.dfg.de/resource/blob/289674/ff57cf46c5ca109cb18533b21fba49bd/230921-stellungnahme-praesidium-ki-ai-data.pdf>

- Gehring, L., & Paaßen, B. (2025). Assessing LLM Text Detection in Educational Contexts: Does Human Contribution Affect Detection? arXiv. <https://doi.org/10.48550/ARXIV.2508.08096>
- Hoffmann, N., & Schmidt, S. (2026). Der Wert des Schreibens. Forschung & Lehre. <https://www.forschung-und-lehre.de/lehre/der-wert-des-schreibens-7450>
- Hüsch, M., Horstmann, N., & Breiter, A. (2026). Künstliche Intelligenz im Studium – die Sicht von Studierenden im Wintersemester 2024/25. DatenCHECK, (06). <https://hochschuldaten.ch.de/kuenstliche-intelligenz-im-studium-die-sicht-von-studierenden-im-wintersemester-2024-25/>
- Lehmann, M., Cornelius, P. B., & Sting, F. J. (2025). AI Meets the Classroom: When Do Large Language Models Harm Learning? (arXiv:2409.09047). arXiv. <https://doi.org/10.48550/arXiv.2409.09047>
- Münzer, S., Go, S., & Paaßen, B. (2026). Gefährdet generative KI die digitale Autonomie im Studium? Diskussionspapier Nr. 41 / Juli 2026. Hochschulforum Digitalisierung. https://hochschulforumdigitalisierung.de/wp-content/uploads/2026/07/HFD_DP_41_Autonomie.pdf
- Perry, N., Srivastava, M., Kumar, D., & Boneh, D. (2023). Do Users Write More Insecure Code with AI Assistants? Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, 2785–2799. <https://doi.org/10.1145/3576915.3623157>
- Wissenschaftsrat (2026). Intellektuelle Souveränität: Empfehlungen für die Hochschulbildung in Zeiten von generativer KI (Drs. 3319-26), Juli 2026. <https://www.wissenschaftsrat.de/download/2026/3319-26>