From "AI is Supi" to "Democracy is in Danger": Users' Perceptions About AI

Dagmar Monett
Computer Science Dept., Berlin School of Economics and Law, Germany

dagmar.monett-diaz@hwr-berlin.de

Abstract: Educational institutions worldwide have introduced courses, training, guidelines, and similar documentation which range from allowing unlimited use of artificial intelligence (AI)-based products to their complete ban in specific settings. Tangentially, and with the spread of chatbots and similar software programs for fast producing textual and other types of information, there are also increasing concerns regarding how and what students actually learn or should, as well as whether outsourcing reading comprehension, academic writing, critical thinking, and other cognitive skills, once paramount in educational contexts, are enhanced or atrophied with the help of AI, and how academic integrity might be devalued as a result of the latter. Tangentially, because the harms are often acknowledged but the use of AI is simultaneously promoted, which is known as responsible realism and is a form of ethics washing. This paper presents part of an empirical, exploratory study whose goals include the analysis of users' perceptions about the use of AI and their awareness of possible harms. To gather opinions on these and related topics, a survey was conducted to which a total of 535 university members participated. All responses were scrutinized using qualitative content analysis. This allowed for the extraction of categories that depict the whole spectrum of their stances regarding AI. The results contribute to understanding users' awareness, opinions, concerns, and needs regarding the use of this technology. They also help in uncovering common tropes that might be preventing university members from assessing generative AI harms critically.

Keywords: AI, AI Harms, Education, Critical AI, Responsible Realism

1. Introduction

Four months after ChatGPT's release on November 30, 2022, an internal webinar on AI-supported writing was organised by the university's IT department to which about 40 people, mainly lecturers, participated. The first presenter opened the webinar with the question *"Will you use ChatGPT in your teaching in Summer 2023?"* The answers from the participants were: *Yes*, 32%; *No*, 25%; and *Not sure*, 42%. The contents then included a brief description about ChatGPT (and generative AI), what could be done with it (at the time), the advantages of using it ("for saving time" being one of them), the legal issues if used for grading students' works, as well as the supposed benefits of "the AI" everybody was talking about. It followed a discussion with a Q&A session. I remember my talking about the AI hype, plagiarism, and issues the technology already had, and how even publishing houses and universities were sharing editorials on either banning its use or clarifying that the only authors should be humans and not software programs. I shared several publications on the chat, like e.g. (Brainard, 2023; COPE, 2023; Flanagin et al., 2023; Nature, 2023; Thorp, 2023). The same question was asked at the end. The answers this time, first to my surprise and then to my disappointment, were: *Yes*, 65%; *No*, 20%; and *Not sure*, 15%, an evident shift that could suggest a lack of awareness of the harms and issues of the technology. How about other university members?

This paper presents part of an exploratory empirical study which goals include, among others, the analysis of university members' awareness of the harms of and their perceptions about AI. It follows the methodology for qualitative content analysis (QCA) suggested in (Mayring, 2014; 2025), which first step consists in *defining concrete research questions* (Mayring, 2014, p. 10). The research questions guiding this work are:

> RQ1: To what extent are university members aware of selected issues and current or potential harms associated with generative AI?
>
> RQ2: What positive and negative perceptions of generative AI are expressed by university members, and to what extent do these perceptions reflect recurring tropes identified in the literature?

The second step of QCA consists in *linking the research question(s) to theory* (Mayring, 2014, p. 11). Much has been written regarding users' perceptions about the use of AI. For instance, Miranda (2025) surveyed university members from an Italian university and found that although positive towards AI, participants were sceptical of its utility for younger learners. Unterweger et al. (2026) document students' concerns about data privacy and academic integrity, among others. In Colbran et al.'s (2026) study about chatbots use, students manifested concerns e.g. about academic integrity and the accuracy of results. Erkılıç and Cifci (2025), in contrast, queried ChatGPT about similar concerns and other academic violations, and suggest that knowing about ChatGPT's issues helps to address problems of unethical behaviour.

Little has changed in some contexts despite abundant research exposing the dangers of the uncritical adoption of generative AI, especially in education (Atelier Ecopol, 2025; Birhane, 2021; Crawford, 2021; Guest et al., 2026; Markelius et al., 2024; Monett & Paquet, 2025). Some users find chatbots' writing "abilities" a perfect substitute for doing the tasks themselves. Others, surrounded by the current wave of AI-based applications, might actually conform to what goes with the AI hype, see technical solutionism as unavoidable, neglect critical harms, and take cognitive growth for granted (Monett & Paquet, 2025).

The next sections introduce the methodological approach underneath QCA and how the data was gathered, analysed, presented and evaluated. I.e. the steps three to seven suggested by Mayring: *definition of the research design*, *definition of the sample or material and the sampling strategy*, *methods of data collection and analysis*, *processing and presentation of results*, and *discussion* (Mayring, 2014, p. 11ff).

2. Methods

The methodological approach used in this research is QCA, "a mixed methods approach: assignment of categories to text as qualitative step, working through many text passages and analysis of frequencies of categories as quantitative step" (Mayring, 2014, p. 10). The text material consists of open-ended answers provided by respondents to survey questions. The research design is *explorative*, where the content analysis is performed following an *inductive* category formation procedure (Mayring, 2014, p. 79ff). The selection criterion includes, but is not limited to, the reasons for or against the use of generative AI in the context of education and research. *Systematically*, categories and sub-categories are built depending on those reasons, i.e. by analysing their *semantic connotation*. General comments, suggestions, concerns, and common tropes (or clichés) are grouped according to their *type*. Their occurrences are quantitatively analysed, *exemplarily*.

To gather university members' opinions, a survey was conducted in July 2025. This paper focuses *only* on the results concerning both participants' perceptions and knowledge about the role that generative AI-based programs are playing in 14 concrete forms of harm as documented in the literature. Reality shows how negative aspects and AI harms are underestimated or ignored not only at the individual but also at the institutional level and beyond. Is educating university members about such harms necessary? How can this knowledge influence institutional policies, teaching, and learning?

2.1 Participants

17855 people were invited to participate in the survey: 13952 students (undergraduate and postgraduate, the latter being e.g. Master's students), 3520 adjunct lecturers (those with teaching contracts and external to the university), 288 professors (those in university positions), and 95 researchers (those being e.g. doctoral students and/or in research positions). The invitation was sent via email to four corresponding university mailing lists which were made available by the IT department of the university. The survey was active two weeks, from July 9 to July 24, 2025, and was provided in two languages, German and English.

A total of 535 people (3%) completed the survey. This should be more than enough for an exploratory study in qualitative analysis (Braun & Clarke, 2025). One respondent belongs to a group not explicitly targeted with the survey and is thus discarded. All other responses (i.e. 534) are analysed, i.e. the sample for QCA is not such but the whole set of responses: 325 from undergraduate (60.9%) and 88 postgraduate (16.5%) students, 104 from adjunct lecturers or professors (19,5%), and 17 from researchers including PhD students (3.2%).

2.2 Data Collection

The survey was created using Google Forms. Options to collect personal data there were deactivated. Apart from questions related to the use of generative AI in teaching and learning, two general questions are key to the analysis presented in this paper: a multiple-choice question comprising 14 current AI harms and one open-ended question to collect any other opinion about generative AI (see Table 1).

2.3 Data Analysis

No generative AI-based tools were used in this research project nor in the production of this paper, the main being a *conscientious objection* (Atelier Ecopol, 2025) to generative AI because of the many proven harms of this technology, especially in, but not limited to, education and academia (see e.g. Guest et al., 2026). Furthermore, Mayring (2025) found that summaries generated by ChatGPT are not suited for QCA and are error-prone,

superficial, and do not recognize hidden text content. Other authors do not recommend the use of generative AI for qualitative analysis either (Jowsey et al., 2025; Nguyen & Welch, 2025; Radu & Rocher, 2026).

Table 1: Questions regarding participants' awareness of AI harms and added opinions on generative AI-related topics (In the German version of the questionnaire, those were questions Q16 and Q17, respectively)

[Q34] Are you aware of the possible role the usage of generative AI plays in... *Check all that apply.* ☐ ...increasing CO2 emissions? ☐...increasing water consumption? ☐...increasing energy consumption? ☐...becoming dependent on AI? ☐...replacing jobs? ☐...extending bias and discrimination? ☐...the loss of knowledge quality? ☐...expanding security breaches? ☐...cognitive offloading? ☐...the loss of human agency? ☐...degrading learning abilities? ☐...violating intellectual property rights? ☐...spreading fake news, misinformation and disinformation? ☐...allowing data workers exploitation? ☐ I am not aware of any of the issues mentioned above.
[Q35] Are there any other comments you would like to add regarding the use of generative AI? Please, add them below.

The results of the open-ended question, Q35 from Table 1 and its corresponding one in German, were analysed manually by following the methodological approach introduced above. Tools for interactive qualitative content analysis (like QCAmap (Mayring, 2025)) were dismissed since no consent from respondents was collected in this regard. In addition, tropes regarding generative AI were analysed by drawing on the notion of *responsible realism*, an instance of ethics washing introduced in (Guersenzvaig & Monett, 2026).

2.4 Quality Criteria

Although carried out carefully and based on the long experience (also in writing) of the author in the topics mentioned above, QCA was done by only one person, i.e., no intra/-inter-coder agreement can be calculated. This is not an issue in qualitative analysis, though (Braun & Clarke, 2025). The results presented here should be considered *exploratory*. To partially counter this limitation, and as suggested by (Mayring, 2014, p. 114), part of the textual material and its categorisation was sent to a colleague, who verified the semantic validity and analysis regarding tropes about generative AI.

3. Results and Discussion

All participants (n=534) answered question Q34 (see Table 1): 479 (89.7%) answered the German and 55 (10.3%) the English version of the questionnaire. Figure 1 shows the total number of occurrences for each harm option, and Table 2 shows their distribution depending on the participant group.

Only 6 participants (1,1% from the total) indicated not being aware of any of the harms. 12 lecturers (11,5% of that group) and 1 researcher (5,9% of that group) do not use generative AI; all of them are aware of at least 2, and 7 of them of at least 11 from 14 harms where generative AI might play a role in. One lecturer indicated not having used generative AI, but in a different question indicated using it for generating ideas and suggestions, which is a contradiction and thus is not counted as not using AI. Several harms rank low for all three groups. As the data shows, there are types of harms where even students know much better than lecturers and researchers. Lecturers, as a group and on average, are actually the least aware of the harms, at least as indicated by the answers they provided. Also possible: some participants might be ignoring the harms or do not see them as

critical enough to stop using AI. This might be a reason for the uncritical adoption of generative AI in their teaching, a serious problem according to (Guest et al., 2026).

96 respondents (18%) answered the open-ended question Q35, 65 of them being students and the rest, lecturers and researchers. 80 comments were collected from the version in German of the questionnaire, and 16 from the one in English. The length of the comments varies from 1 to 327 words.

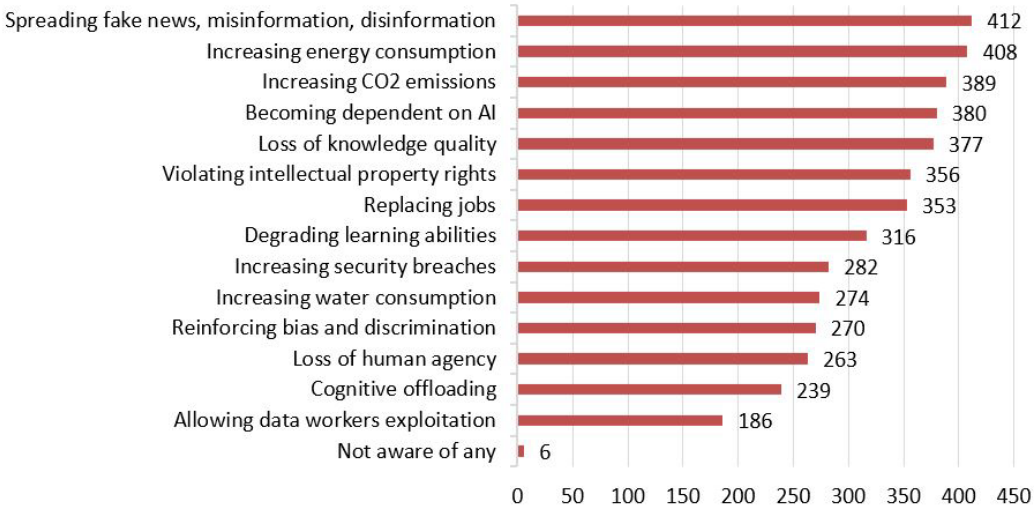


Figure 1: Participants' acknowledged awareness: Possible role the usage of generative AI plays in different harms

Table 2: Participants' acknowledged awareness: Distribution and percentage relative to the participant group (emphasis: top 3 highest and top 3 lowest scores)

Awareness of…	Students	Lecturers	Researchers
Spreading fake news, misinformation, disinformation	325 (**78,7%**)	73 (70,2%)	14 (**82,4%**)
Increasing energy consumption	312 (**75,5%**)	82 (**78,8%**)	14 (**82,4%**)
Increasing CO2 emissions	294 (71,2%)	81 (**77,9%**)	14 (**82,4%**)
Becoming dependent on AI	306 (**74,1%**)	61 (58,7%)	13 (76,5%)
Loss of knowledge quality	295 (71,4%)	72 (69,2%)	10 (58,8%)
Violating intellectual property rights	266 (64,4%)	78 (**75,0%**)	12 (70,6%)
Replacing jobs	285 (69,0%)	58 (55,8%)	10 (58,8%)
Degrading learning abilities	238 (57,6%)	68 (65,4%)	10 (58,8%)
Increasing security breaches	216 (52,3%)	55 (52,9%)	11 (64,7%)
Increasing water consumption	217 (52,5%)	44 (**42,3%**)	13 (76,5%)
Reinforcing bias and discrimination	188 (**45,5%**)	69 (66,3%)	13 (76,5%)
Loss of human agency	219 (53,0%)	36 (**34,6%**)	8 (**47,1%**)
Cognitive offloading	187 (**45,3%**)	45 (43,3%)	7 (**41,2%**)
Allowing data workers exploitation	143 (**34,6%**)	35 (**33,7%**)	8 (**47,1%**)
Not aware of any	4 (1,0%)(*)	1 (1,0%)(**)	1 (5,9%)

*: 3 other students also indicated being aware of almost all the harms, which is a contradiction.

**: 1 other researcher also indicated being aware of almost all the harms, which is a contradiction.

One by one, each comment was analysed, thoroughly, and its semantic connotation categorised. The following broad sets were constructed along this process, inductively (concrete examples are showed in Tables 3 to 5):

- General comments (positive or negative) regarding generative AI (e.g. related to the respondent's experience)

- Concerns regarding the use of generative AI (i.e. concrete concerns the respondent has)
- Needs or suggestions mentioned by the respondent (some respondents mention them explicitly)

The categories and sub-categories, created also inductively from the topics addressed in the comments, were:

- Education (assessment, grading, curriculum, teaching, training, learning, learning methods, cognitive abilities, cognitive skills, human abilities)
- Ethics (awareness, ethical aspects, good practice, requirement)
- AI use (AI in teaching, AI in learning, AI in research, software coding)
- Issues (AI limitations, challenges, harms, risks, laziness, lock-in, stolen data, plagiarism, replacement, jobs)
- Governance (guideline, regulation, banning)
- Pushback (refusal, opposition, denial)
- Tropes (AI hype, fear of missing out (FOMO), ignorance about functioning, inevitability, category error, false generalization, uncritical adoption, responsible realism, industry push, efficiency, time)
- Survey (content-related)

These categories did not exist *a priori* but were built *systematically* after analysing each of the 96 comments that were provided. From those categories, the most frequently mentioned topics were *AI use* (e.g. whether AI should be used or not, as a support tool or not, good in research or not, good for learning or not, etc.) with more than 60 mentions in total; *harms* (e.g. to individuals, study programme, institution, society, etc.) with more than 50 mentions, and *tropes* (see below) with more than 45 mentions.

Table 3: Examples of general aspects regarding generative AI, grouped by semantic connotation

Positive Aspects:	Negative or Problematic Aspects:
Increasing importance of AI-related skills	AI is promoted everywhere, especially at the university
Use of AI is important in education	AI is in all products without being asked; it interrupts work
AI can be good in research	There is an overuse of AI
AI is a very important tool for learning	AI is not only ChatGPT
AI is helpful for coding	Regulation does not prevent use
AI should be freely available for use in writing	There is no motivation not to use AI
AI for spell and grammar check is Ok	AI as a tool is good for those who know how to check it
AI is useful	Not all students use AI but there is pressure from the university
AI is supi	Lecturers using AI is not well seen
	Lesser suitable assessment was needed because of AI
	AI is pushed by industry
Both positive and negative aspects in the same comment:	Other general aspects:
Generative AI has both pros and cons	Developing own's cognitive abilities and experiences is important for learning and in education
Use in research can be positive but also problematic	Own's achievements and learning shouldn't be replaced (by tools' output)
AI makes the work easier but not always better	Students who use AI have advantage but learn less
	There are law suits against AI

Table 4: Examples of participants' concerns regarding the use of generative AI, grouped by semantic connotation

Related to Teaching and Learning:	Involving Duty, Responsibility, or Ethical Behaviour:
There is no possibility to know about AI usage	There is no space for opposition (to the use of AI)
The learning effect when using AI is smaller or poorer	There is disregard for harms
The use of AI foments laziness	Licenses are purchased despite the harms or without consulting users
Students replace thinking completely with AI	

Related to Teaching and Learning:	Involving Duty, Responsibility, or Ethical Behaviour:
Some students use it to produce whole writing works	Adoption without consulting; technology lock-in
Group work is affected (when some use AI and others not)	Students are not aware of or clarified about their responsibility for the results
Students' use of AI should be disclosed but who will check anyways?	There is little awareness about responsibility structures
Lecturers often fail to explain content well; thus, AI should be used	There is no discussion about ethical considerations and negative impact
There is a possible skill degradation	Banning means disadvantage
There might be unfairness when grading	There is AI hype
Use of AI should not be supported at the university	Some of the effects are controversial or have not been clearly proven
Grade should not be affected if students use AI because lecturers cannot explain content well	Used wrong, one could learn little, become dependent, and be replaceable
Plagiarism is an issue	Replacing own's actions completely is problematic
Accusing someone of plagiarism based on the use of AI is unjustified	
Excessive and especially undeclared use of AI lowers the academic standards of the program	
Related to the Technology:	**Other General Concerns:**
The output of generative AI is often false or not existent	Not using AI is seen as a “blocker for progress”
Use should be careful because of errors or non-existent sources	Not using AI is a career disadvantage
Output should not be trusted; it should be verified first	Tasks that once made fun might disappear
Output does not meet scientific standards	AI exacerbates harms
That using it saves time is questionable	There are many harms and risks
Many people have no knowledge about how AI functions	Interests of a few over interests of humanity
AI has limitations like dependency on training (stolen) data	There is a potential for threats to life
Online content is flooded with AI-generated content	

Table 5: Needs or suggestions, grouped by semantic connotation

Related to Teaching and Learning:	Involving Duty, Responsibility, or Ethical Behaviour:
Session on how to use AI efficiently in research and teaching	Use of AI must be disclosed
Faculty must be trained to deal with ChatGPT	Use of AI should be allowed
How to let students use AI	Use of AI should be regulated and controlled by experts
How to test/verify the use of AI when grading	AI should not be banned, but regulated
How to use AI to personalize asynchronous learning	AI output must be assessed critically
New assessment forms are needed	Responsibility should not be outsourced to faculty
A basic understanding of software coding is needed	AI tools should be used with moderation
Teacher training on the good and also on the risks of AI	Pros and cons of the use of AI should be evaluated
More information about pros and cons for learning	Discussion on ethical implications is needed
Learning activities must be adapted, cannot be the same	Faculty should be provided with clear guidelines
AI should be integrated into the learning process	Guidelines not only for students, but also for faculty
AI should be freely available for use in writing	
Related to the Technology:	**Other General Suggestions:**
More training of AI on human research methods	Need to determine which human tasks make sense or not anymore
Prompt engineering as lecture	Differentiating between the human and the artificial will be essential
AI compliance with data protection regulations when doing research is needed	

The following examples of comments contain instances of the category *AI use*:

> *"AI as support is good as long as it is not used to write complete works."* (Lecturer; translated from German)

> *"I hate AI, but sometimes by deadlines and academic stress means that I have no choice but to save time in some instances by finding literature and clarifying concepts. However, when I have done this in the past ChatGTP had produced information that is not factually correct, so I don't actually know if it saves me time."* (Postgraduate student)

The following examples of comments contain instances of the category *harms*:

> *"There is a huge disregard for the harms of generative AI."* (Lecturer)

> *"The welfare of humanity must take precedence over the interests of a few. Unfortunately, that is not the case at present. AI is exacerbating the problems caused by digitalization.* ***Democracy is in danger."*** (Undergraduate student; translated from German)

Reflecting on the users' perceptions about AI, and as manifested in the comments to the open-ended question Q35, it is evident that there is much to be learnt and done regarding generative AI. The whole spectrum of opinions—from positive to negative, for or against AI, from demonstrably true to demonstrably false, from ethical to unethical—is present across all participant groups. Some perspectives, however, ignore the limits of AI (Guest et al., 2026; Selwyn, 2024).

Drawing on the concept of *responsible realism*, a form of ethics washing and "stance in which Higher Education institutions [or individuals] acknowledge the harms and ethical problems of commercial AI, yet still promote its use, shifting the burden of mitigation onto individual end users by calling to 'responsible use'" (Guersenzvaig & Monett, 2026, p. 1), each comment was analysed to detect if it contained traces of this kind of narrative or similar ones. As expected, responsible realism was found in the respondents' comments, at least in 13 of them (14,8%). Several examples are presented in Table 6, the last column being their type.

Table 6: Examples of common tropes (fallacies, clichés) about generative AI, together with their type

Trope, Fallacy, Cliché	Type
Assuming that chatting with a software program is an important skill	Overvalue of skill
No AI use is a blocker for progress	Techno-solutionism
AI cannot be stopped	Inevitability
AI is useful and not problematic when it functions right and someone is responsible	Ethical washing
The tool is good if the output is analysed critically	Critical washing
AI will dramatically transform higher education and research	AI hype, inevitability
"Prompt instructor" will be a job opportunity	AI hype
Banning generative AI makes no sense	Devaluing of regulation
Education should be transformed to the "new reality" of AI	Determinism
AI is increasingly relevant for companies; thus, it should be taught at the university	Inevitability, category error
Those using AI will replace those not using it	FOMO
Potential or competitiveness not possible if students do not use AI or learn how to use it	FOMO
AI is just a tool	Disregard for biases
AI is a partner	Anthropomorphism
Everybody uses AI; those who say they don't are lying	False generalisation
Teach how to use it properly and issues will be mitigated	Wishful thinking, critical washing
Delegating cognitive tasks to AI is good: more time available and fun when learning	Ignorance about how learning occurs
Harms depend on how AI is used	Ignorance about how the technology works, ethical washing

There was no category formation in this part of the analysis. Those are well-known tropes, easy to detect based on the author's experience and publications, and as defined in the research literature (see e.g. Guersenzvaig & Monett, 2026; Guest et al., 2026; Morozov, 2013; Postman, 1992).

Based on determinism and inevitability, and recreating cases of techno-solutionism, anthropomorphism, AI hype, and even ignorance about both how generative AI actually works and the socio-technical conditions that determine it, responsible realism provides a cover for unethical practices, unfortunately. Those users' perspectives might risk undermining the goals of education (Postman, 1995). As Selwyn (2024) reminds us:

> "The future of AI and education is *not* a foregone conclusion that we simply need to adapt to. Instead, the incursion of AI into education is definitely something that can be resisted and reimagined."

4. Conclusion

The findings presented above provide a varied perspective regarding the research questions. Yet, they should be taken as a casual snapshot of much more complex reasons, dynamics, and conditions regarding the use of generative AI. People, tools, and users' perceptions are constantly changing. Refusal, rejection, and organisation are evolving too. The findings were considered when designing two micro-credentials courses on Critical AI that will be offered to university members starting 2027. Interested readers may reach out to request more details about them.

Other university members were excluded from participating, but staff and other administrative members are also generative AI users and might influence its adoption. Another limitation is the possible bias some respondents might have had when answering the open-ended question: at least seven of them were visibly upset with the question about their awareness of the harms. It is known, for example, that the rejection of these technologies and criticism of the *status quo* is often accompanied by erasure or even retaliation (Täuber, 2026). To partially prevent this, anonymous participation was considered. But if there is no open culture for discussion about the most acute AI issues affecting education and research, we cannot expect sincerity and a civilised discussion, needless to mention principled action.

Acknowledgements

Especial thanks go to Ariel Guersenzvaig for the fruitful discussions on responsible realism, to Maja Günther and Titus Reiche for their support in preparing the questionnaire, and to the anonymous participants to the survey.

Ethics Declaration

No personal data was collected. The questionnaire was reviewed by the central division "Research Support Services" of the university and officially approved on July 7th, 2025.

AI Declaration

No generative AI was used. Translations were done by the author.

References

Atelier Ecopol (2025, December 4) "Faced with generative AI, conscientious objection" (Original title in French: Face à l'IA générative, l'objection de conscience), *ATelier d'ÉCOlogie POLitique*, https://informatique.up8.edu/actu/2025-2026.html#2025-12-04_17-30

Birhane, A. (2021) "The Impossibility of Automating Ambiguity", *Artificial Life*, 27(1):44–61.

Brainard, J. (2023) "As scientists explore AI-written text, journals hammer out policies", *Science*, *379*(6634):740–741.

Braun, V., and Clarke, V. (2025) "Reporting guidelines for qualitative research: a values-based approach", *Qualitative Research in Psychology*, *22*(2):399–438.

Colbran, S., Jha, M., and Schiavone, C. (2026) "Understanding student perspectives on generative AI chatbots: a human-centred mixed-methods study in higher education", *International Journal of Educational Technology in Higher Education*, *23*, 28.

COPE (2023) "Authorship and AI tools", *COPE Council*, Committee on Publication Ethics, https://doi.org/10.24318/cCVRZBms

Crawford, K. (2021) *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*, Yale University Press.

Erkılıç, E., and Cifci, I. (2025) “Ethical Issues and Violations in Using Chatbots in Academic Writing and Publishing: The Answers from ChatGPT”, *Journal of Multidisciplinary Academic Tourism*, *10*(1):111–120.

Flanagin, A., Bibbins-Domingo, K., Berkwits, M. and Christiansen, S.L. (2023) “Nonhuman ‘Authors’ and Implications for the Integrity of Scientific Publication and Medical Knowledge”, *Journal of the American Medical Association*, *329*(8):637–639.

Guest, O., Suárez, M., Müller, B.C.N. and van Meerkerk, E. et al. (2026) “Against the Uncritical Adoption of 'AI' Technologies in Academia”, *Digital Culture & Education*, *16*(2):85–118.

Jowsey, T., Braun, V., Clarke, V., Lupton, D. and Fine, M. (2025) “We Reject the Use of Generative Artificial Intelligence for Reflexive Qualitative Research”, *Qualitative Inquiry*, https://doi.org/10.1177/10778004251401851

Markelius, A., Wright, C., Kuiper, J., Delille, N. and Kuo, Y.-T. (2024) “The mechanisms of AI hype and its planetary and social costs”, *AI Ethics*, *4*:727–742.

Mayring, P. (2014) *Qualitative Content Analysis: Theoretical foundation, basic procedures and software solution*, Social Science Open Access Repository, https://www.ssoar.info/ssoar/handle/document/39517

Mayring, P. (2025) “Qualitative Content Analysis With ChatGPT: Pitfalls, Rough Approximations and Gross Errors. A Field Report”, *Forum Qualitative Social Research*, *26*(1).

Miranda, S. (2025) “Artificial Intelligence in Education: An Exploratory Survey to Gather the Perceptions of Teachers, Students, and Educators Around the University of Salerno”, *Education Sciences*, *15*(8):975.

Monett, D. and Paquet, G. (2025) “Against the commodification of education — if harms then not AI”, *Journal of Open, Distance, and Digital Education*, *2*(1):1–24.

Morozov, E. (2013) *To save everything, click here: The folly of technological solutionism*, PublicAffairs.

Nature (2023) “Tools such as ChatGPT threaten transparent science; here are our ground rules for their use”, *Nature*, *613*:612.

Nguyen, D.C. and Welch, C. (2025) “Generative Artificial Intelligence in Qualitative Data Analysis: Analyzing—Or Just Chatting?” *Organizational Research Methods*, *29*(1):3–39.

Postman, N. (1992) *Technopoly: The surrender of culture to technology*, Vintage Books.

Postman, N. (1995) *The end of education: Redefining the new value of school*, Vintage Books.

Radu, R. and Rocher, L. (2026) “Attributing and situating knowledge cannot be left to language models”, *Nature Machine Intelligence*, *8*:295.

Selwyn, N. (2024) “On the Limits of Artificial Intelligence (AI) in Education”, *Nordisk tidsskrift for pedagogikk og kritikk*, *10*(1):3–14.

Täuber, S. (2026) “Erasing critical scholarship: How attacks on academic freedom undermine democracy”, *Organization*, 1–19.

Thorp, H.H. (2023) “ChatGPT is fun, but not an author”, *Science*, *379*(6630):313.

Unterweger, M., Hörmann, C., Kuka, L., Sabitzer, B. (2026) “The AI generation at school: A large-scale survey on student practices, perceptions, and worries”, *Social Sciences & Humanities Open*, *13*, 102821.