


**UNIVERSITY OF
LUCERNE**

**FACULTY OF HUMANITIES AND
SOCIAL SCIENCES**

SOCIOLOGICAL SEMINAR

WORKING PAPER



BETWEEN ADOPTION AND DOUBT

**A CENSUS OF MORAL ATTITUDES TOWARDS ARTIFICIAL INTELLIGENCE AMONG
STUDENTS AT THE UNIVERSITY OF LUCERNE**

PATRICK SCHENK | VANESSA MÜLLER | LUCA KEISER

Patrick Schenk, Vanessa Müller, Luca Keiser

Between Adoption and Doubt: A Census of Moral Attitudes towards Artificial Intelligence among Students at the University of Lucerne

Working Paper

January 2024

University of Lucerne, Switzerland

DOI 10.5281/zenodo.10405081

Author information:

Patrick Schenk is a postdoctoral researcher at the University of Lucerne, Switzerland. The main focus of his research is economic sociology and consumption, cultural sociology, the sociology of morality, and the sociology of technology.

Vanessa Müller is a PhD candidate in sociology at the University of Lucerne, Switzerland. In her research she focuses on finding out how the societal implementation of artificial intelligence changes the concept of being human.

Luca Keiser is a political data scientist at the gfs.bern, Switzerland. His main focus is on preparing and analyzing data of all kinds, as well as the (visual) communication of the results.

Corresponding author:

Patrick Schenk

Department of Sociology, University of Lucerne

Frohburgstrasse 3, 6002 Lucerne, Switzerland

E-mail: patrick.schenk@unilu.ch

Acknowledgements:

Special thanks are due to Gabriel Abend for his invaluable contributions to this project and his comments, Lukas Portmann for making the census at the University of Lucerne possible, Cornelia Sidler for her ongoing support, and finally Giulia Pellanda and Linda Widmer for their impeccable translations.

Between Adoption and Doubt: A Census of Moral Attitudes towards Artificial Intelligence among Students at the University of Lucerne © 2024 by Patrick Schenk, Vanessa Müller, Luca Keiser is licensed under CC BY-NC 4.0. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/4.0/>

Cover image generated on 20.12.2023 by Open AI's DALL-E in ChatGPT4 with the prompts: 1: Generate an image suitable for the cover of a scientific report with the title: Between Adoption and Doubt: A Census of Moral Attitudes towards Artificial Intelligence among Students at the University of Lucerne. 2: Generate an image which has to do with artificial intelligence and university students but doesn't have actual human beings on it. and edited by Patrick Schenk and Daniel Jurt. The cover image is not licensed under CC-BY-NC 4.0

Executive summary

1) As early adopters and key carriers of social activism and innovation, university students are a particularly important group for understanding the social ramifications of AI-based technologies. To get a more accurate picture of their use, perceptions, and moral evaluations of AI, we conducted a full census of students at the University of Lucerne (Switzerland). From May to June of 2023, 440 students from all faculties participated in a standardized online survey.

2) Students are highly interested users of AI-based technologies. 98 percent of respondents have experiences with intelligent translators and 70 percent with chatbots, including the popular example of ChatGPT. Respondents are primarily interested in the moral consequences of new technologies, followed by legal frameworks, personal use, and economic consequences.

3) Despite the widespread use and interest, we observe high levels of uncertainty. The majority of respondents is undecided whether to trust AI or not. Additionally, only about every fifth respondent feels confident about their knowledge of AI. Thus, when it comes to AI-based technologies, students do not correspond to the widely employed figure of a digital native.

4) In general, students are highly skeptical of the future impact of artificial intelligence on society. Around 60 percent believe AI is going to take away jobs and reduce privacy. In contrast, only 22 percent believe AI will improve society. The largest disagreement concerns the question whether advanced AI poses a threat to the self-esteem of human beings.

5) A substantial number of students would still be willing to include AI in solving moral and political problems. One in four students thinks that AI is more likely than the Swiss population to be able to recognize optimal political decisions. Every fifth student would ask AI for moral advice. At the same time, students are reluctant to attribute moral responsibility, moral worth, or moral rights to AI. This raises new questions for the distribution of moral agency and accountability among human and artificial agents.

6) We derive various practical implications from these findings. Universities should enable students to assess the trustworthiness of AI and foster skills to understand and explain these new technologies. They should provide spaces to discuss the morality of AI, especially employment effects, environmental impacts, and human self-esteem, while avoiding extreme predictions and alarmist accounts of the future of AI, whether dystopian or utopian.

Zusammenfassung

1) Als «Early-Adopters» und Träger:innen von Innovation und sozialem Aktivismus sind Universitätsstudierende eine besonders wichtige Gruppe, um die sozialen Auswirkungen von KI-basierten Technologien zu verstehen. Um ein genaueres Bild ihrer Nutzung, Wahrnehmung und moralischen Bewertung von KI zu erhalten, führten wir eine Vollerhebung unter Studierenden der Universität Luzern (Schweiz) durch. Von Mai bis Juni 2023 beteiligten sich 440 Studierende aus allen Fakultäten an einer standardisierten Online-Befragung.

2) Studierende sind sehr interessierte Nutzer:innen von KI-basierten Technologien. 98 Prozent der Befragten haben bereits Erfahrungen mit intelligenten Übersetzungsprogrammen und 70 Prozent mit Chatbots (z.B. ChatGPT) gesammelt. Die Befragten interessieren sich insbesondere für die moralischen Folgen neuer Technologien, gefolgt von den rechtlichen Rahmenbedingungen, der persönlichen Nutzung und den wirtschaftlichen Folgen.

3) Trotz der verbreiteten Nutzung und dem grossen Interesse beobachten wir eine ausgeprägte Unsicherheit. Die Mehrheit der Befragten ist unentschlossen, ob sie KI vertrauen soll oder nicht. Ausserdem betrachten nur etwa 20 Prozent der Befragten ihr eigenes Wissen zu KI als ausreichend. Wenn es um KI-basierte Technologien geht, entsprechen Studierende nicht dem weit verbreiteten Bild des «Digital Natives».

4) Generell sind die Studierenden sehr skeptisch, was die zukünftigen Auswirkungen künstlicher Intelligenz auf die Gesellschaft betrifft. Rund 60 Prozent glauben, dass KI Arbeitsplätze vernichten und die Privatsphäre einschränken wird. Im Gegensatz dazu denken nur 22 Prozent, dass KI die Gesellschaft besser machen wird. Die grösste Uneinigkeit besteht bei der Frage, ob fortgeschrittene KI eine Bedrohung für den Selbstwert des Menschen darstellt.

5) Eine beträchtliche Anzahl von Studierenden wäre dennoch bereit, KI in die Lösung moralischer und politischer Probleme einzubeziehen. Ein Viertel der Studierenden glaubt, dass KI bessere politische Entscheidungen als die Schweizer Bevölkerung treffen könnte. 20 Prozent würden KI für moralischen Rat heranziehen. Jedoch zögern die Studierenden, KI moralische Verantwortung, moralischen Wert oder moralische Rechte zuzuschreiben, was neue Fragen zur Verteilung moralischer Handlungsfähigkeit und Verantwortlichkeit zwischen menschlichen und künstlichen Agenten aufwirft.

6) Wir leiten aus diesen Ergebnissen praktische Implikationen ab. Universitäten sollten Studierende befähigen, die Vertrauenswürdigkeit von KI einzuschätzen, und zudem die Kompetenz fördern, diese neuen Technologien zu verstehen und erklären zu können. Sie sollten Diskussionsräume über die ethischen Aspekte der KI, einschliesslich ihrer Auswirkungen auf Arbeit, Umwelt und menschlichen Selbstwert, schaffen, und gleichzeitig extreme Vorhersagen und alarmistische Deutungen der Zukunft mit KI vermeiden.

Résumé exécutif

1) En tant qu'adopteurs précoces et vecteurs clés de l'activisme social et de l'innovation, les étudiants universitaires constituent un groupe particulièrement important pour comprendre les ramifications sociales des technologies basées sur l'IA. Pour obtenir une image plus précise de l'utilisation, des perceptions et des évaluations morales de l'IA, nous avons effectué un recensement complet des étudiants de l'Université de Lucerne (Suisse). De mai à juin 2023, 440 étudiants de toutes les facultés ont participé à une enquête en ligne standardisée.

2) Les étudiants sont des utilisateurs très intéressés par les technologies basées sur l'IA. 98 % des répondants ont déjà utilisé des traducteurs intelligents et 70 % des chatbots, y compris l'exemple populaire de ChatGPT. Les répondants s'intéressent principalement aux conséquences morales des nouvelles technologies, suivies des cadres juridiques, de l'utilisation personnelle et des conséquences économiques.

3) Malgré l'utilisation et l'intérêt généralisés de l'IA, nous observons des niveaux élevés d'incertitude. La majorité des personnes interrogées sont indécises quant à savoir si elles doivent faire confiance à l'IA ou non. En outre, seul un répondant sur cinq environ se sent confiant dans ses connaissances de l'IA. Ainsi, en ce qui concerne les technologies basées sur l'IA, les étudiants ne correspondent pas à la figure largement répandue du natif numérique.

4) En général, les étudiants sont très sceptiques quant à l'impact futur de l'intelligence artificielle sur la société. Environ 60 % d'entre eux pensent que l'IA va supprimer des emplois et réduire la protection de la vie privée. En revanche, seuls 22 % d'entre eux pensent que l'IA améliorera la société. Le désaccord le plus important concerne la question de savoir si l'IA perfectionnée constitue une menace pour l'estime de soi des êtres humains.

5) Un nombre substantiel d'étudiants serait encore prêt à inclure l'IA dans la résolution de problèmes moraux et politiques. Un étudiant sur quatre pense que l'IA est plus à même que la population suisse de reconnaître les décisions politiques optimales. Un étudiant sur cinq demanderait des conseils moraux à l'IA. En même temps, les élèves sont réticents à attribuer à l'IA une responsabilité morale, une valeur morale ou des droits moraux. Cela soulève de nouvelles questions quant à la répartition de l'action morale et de la responsabilité entre les agents humains et artificiels.

6) Nous tirons diverses implications pratiques de ces résultats. Les universités devraient permettre aux étudiants d'évaluer la fiabilité de l'IA et de développer des compétences pour comprendre et expliquer ces nouvelles technologies. Elles devraient offrir des espaces pour discuter de la moralité de l'IA, en particulier des effets sur l'emploi, de l'impact sur l'environnement et de l'estime de soi-même, tout en évitant les prédictions extrêmes et les récits alarmistes sur l'avenir de l'IA, qu'ils soient dystopiques ou utopiques.

Riassunto esecutivo

1) In qualità di “early adopter” e vettori chiave dell’attivismo sociale e dell’innovazione, gli studenti universitari sono un gruppo particolarmente importante per comprendere le ramificazioni sociali delle tecnologie basate sull’intelligenza artificiale. Per ottenere un quadro più accurato del loro utilizzo, delle loro percezioni e delle loro valutazioni morali dell’IA, abbiamo condotto un censimento completo degli studenti dell’Università di Lucerna (Svizzera). Da maggio a giugno del 2023, 440 studenti di tutte le facoltà hanno partecipato a un sondaggio on-line standardizzato.

2) Gli studenti sono utenti molto interessati delle tecnologie basate sull’intelligenza artificiale. Il 98% degli intervistati ha esperienze con traduttori intelligenti e il 70% con chatbot, incluso il popolare esempio di ChatGPT. Gli intervistati sono principalmente interessati alle conseguenze morali delle nuove tecnologie, seguite dai quadri giuridici, dall’uso personale e dalle conseguenze economiche.

3) Nonostante l’uso diffuso e l’interesse, osserviamo alti livelli di incertezza. La maggior parte degli intervistati è indecisa se fidarsi o meno dell’IA. Inoltre, solo un intervistato su cinque si sente sicuro della propria conoscenza dell’IA. Pertanto, quando si tratta di tecnologie basate sull’intelligenza artificiale, gli studenti non corrispondono alla figura ampiamente utilizzata di un nativo digitale.

4) In generale, gli studenti sono molto scettici sull’impatto futuro dell’intelligenza artificiale sulla società. Circa il 60% ritiene che l’IA toglierà posti di lavoro e ridurrà la privacy. Al contrario, solo il 22% crede che l’IA migliorerà la società. Il più grande disaccordo riguarda la questione se l’IA avanzata rappresenti una minaccia per l’autostima degli esseri umani.

5) Un numero considerevole di studenti sarebbe ancora disposto a includere l’IA nella risoluzione di problemi morali e politici. Uno studente su quattro ritiene che l’IA sia più propensa della popolazione svizzera a riconoscere le decisioni politiche ottimali. Uno studente su cinque chiederebbe all’IA un consiglio morale. Allo stesso tempo, gli studenti sono riluttanti ad attribuire responsabilità morale, valore morale o diritti morali all’IA. Ciò solleva nuove questioni per la distribuzione dell’arbitrio morale e della responsabilità tra agenti umani e artificiali.

6) Da questi risultati derivano varie implicazioni pratiche. Le università dovrebbero consentire agli studenti di valutare l’affidabilità dell’IA e promuovere le competenze per comprendere e spiegare queste nuove tecnologie. Dovrebbero fornire spazi per discutere la moralità dell’IA, in particolare gli effetti sull’occupazione, gli impatti ambientali e l’autostima umana, evitando previsioni estreme e resoconti allarmistici sul futuro dell’IA, sia distopici che utopici.

Table of contents

1	Introduction.....	1
2	Use of AI.....	1
3	Knowledge and Interests	2
4	Assessing the Future Consequences of AI	4
5	Moral Evaluation of AI.....	5
6	Moral Rights and Moral Value of AI	7
7	Conclusion.....	8
	Appendix A. Data and Methods	10
	References	11

1 Introduction

With the launch of ChatGPT, general awareness of artificial intelligence (AI) has been growing exponentially. The accelerated integration of AI-based technologies in our society, not only in the form of chatbots, but also in other contexts such as social media, self-driving cars or in social areas such as medicine, journalism or science, raises pertinent questions about morality and knowledge production (Bostrom, 2014; Mittelstadt et al., 2016). With AI taking on more and more tasks previously reserved for human agents, societies are confronted with challenges relating to the quality of non-human generated outputs, the assessment of the future effects of AI, and new moral responsibilities and rights of artificial and human agents.

Universities are a particularly important site for understanding the social ramifications of AI-based technologies. First, students are early adopters of such technologies. Universities are hence forced to deal with the increasing use of AI in higher education (Chan, 2023; study.com, 2023). In addition to concerns over academic integrity and fraud, there is a need to find appropriate ways to avoid over-reliance on AI-based outputs, often generated in an opaque process (Litvinski, 2023). The University of Lucerne is no exception. Events debating questions like “How can universities use AI? And what do they have to keep in mind while using AI?” (Mihálka, 2023) clearly indicate the demand for ethics and well-founded knowledge on this topic. Second, university students are an especially important social group to study the use, perceptions, and attitudes of and towards AI from a social science perspective. Students are key carriers of social activism and innovation in various social fields, from science, to entrepreneurship, to culture and politics (Bernardino & Freitas Santos, 2022; Dahlum & Wig, 2021). They therefore have the potential to play a significant role in shaping the future implementation of AI. Empirical studies on the use and perceptions of AI among university students are hence essential (Almaiah et al., 2022; Gherheş, 2018). They help to identify the benefits, potentials, and problems of AI as they present themselves from the perspective of these social actors. They also lay the ground for developing effective ways for teaching and learning about these technologies (Chan, 2023).

For these reasons, we conducted a full census of the students at the University of Lucerne (Switzerland) from all faculties. We collected data on the use, perceptions, and moral evaluations of AI by means of a standardized online survey. The survey took place in May and June of 2023. 440 students participated (Appendix A describes the methodological procedures more in detail).¹ We focused on three interrelated questions: First, we examined the engagement with AI, including the actual use of AI-based technologies, interest in AI, and knowledge on AI. Second, we surveyed the assessment of various consequences of AI discussed in academic and public debate, such as unemployment and environmental impact. Lastly, we investigated moral evaluations of AI, including trust, moral responsibility, and moral rights. In the following, we first present our empirical findings before we derive practical implications for universities in the concluding section.

2 Use of AI

Applications of AI-based technologies are multiplying. Some of these applications, such as generative AI (e.g., ChatGPT or Midjourney) or self-driving cars, have received increasing media attention. Yet, in the midst of an AI-hype, the pervasiveness of these technologies might be overstated. To gain a more accurate picture, we

¹ The study is part of a larger project funded by the Swiss National Science Foundation on AI and moral decision making, see www.unilu.ch/news/kuenstliche-intelligenz-und-moralische-entscheidungen-5907.

asked students at the University of Lucerne how often they are using various AI-applications. Figure 1 depicts the frequency of using well-known examples of these technologies in decreasing order.




Figure 1: Frequency of using AI-based technologies (n=280-281)

We find that the actual use of these technologies varies greatly. Some of them, namely intelligent translators (e.g., DeepL) and social media, are being used by nearly all respondents, even on a regular basis. Intelligent translators are especially interesting from an academic perspective. To a certain extent, competencies and tasks regarding foreign languages, traditionally part of higher education, are being delegated to AI. For a second group of technologies, we find that a sizeable number of respondents have some experience with them, yet they are only being used by a minority in a regular manner. Chatbots, including the popular example of ChatGPT, belong to this group. While a substantial proportion of 70 percent has at least some experience with chatbots, only 27 percent use them regularly. Around every third student has no experience with chatbots at all. The intensity of using chatbots thus clearly differs between students. Nevertheless, given the host of applications of this new technology in academic and public life, this finding supports the need for a stronger discussion of its implications. The other two technologies belonging to this group are image recognition software and voice assistants. They exhibit a similar pattern, with voice assistants showing a substantially lower proportion of regular users. Looking at the remaining technologies, the prevalence drops drastically. The proportion of students having experience with image generators, autonomous vehicles, or robo advisors is well below 20 percent. This does not necessarily imply a dismissal of AI technologies. It might simply be a consequence of limited general demand for these use cases (e.g., image generators) or lack of opportunities (e.g., autonomous vehicles).

3 Knowledge and Interests

To arrive at a better understanding of the engagement with AI, we also asked questions concerning self-perceived knowledge about and interest in AI. Students perceive their knowledge of artificial intelligence as rather

limited. Less than one third of the respondents agrees with the statement that they know which products and services use AI (Figure 2, “Knowledge Products and Services”). Slightly more students disagree. The number of students indicating that they know a lot about artificial intelligence is even lower with 20 percent (“Knowledge AI”). 40 percent of the students disagree with that statement. Thus, the majority of respondents does not feel particularly competent when it comes to AI. The perceived deficit in AI knowledge among students suggests a need to foster AI literacy.

These findings are surprising given the widespread view that students born post-1980 are considered as digital natives. Being a digital native means a presumed affinity with advanced knowledge and skills in new technologies compared to those born prior to 1980. Following this assumption, authors and public actors demand educational reforms to meet the emerging needs of this new generation of learners (Prensky, 2001). However, our results on the students’ self-perceived proficiency in AI challenges this view. The figure of the digital native might not apply after all, or, at least not in the sense of a natural superiority in the literacy of AI-based technologies.




Figure 2: Knowledge of AI and interest in AI (n=280-282)

Our study also shows a clear majority of students being generally interested in the topic of artificial intelligence (Figure 2, “Interested in AI”). Most telling, only around every fifth person disagrees with the statement that they are interested in AI. Additionally, we collected more fine-grained data on specific topics related to new technologies. Among eight topics, respondents were able to indicate which topics are of particular interest to them. Looking first at the social dimension of new technologies, depicted on the left-hand side of Figure 3, we see that moral concerns rank first, followed by legal frameworks and economic consequences. At least half of the students selected each of these topics, with moral consequences being selected by a notable 68 percent. The history of modern technologies is only of interest to a smaller subgroup. Turning to topics more closely related to the technological artefacts themselves, depicted on the right-hand side of Figure 3, we find that instrumental aspects rank first (personal use and functionality), while prices, availability, and the design of technologies are of secondary interest. Taken together, these results clearly demonstrate the importance of the moral consequences of AI for our respondents. Students are eager to deal with these moral challenges. Hence, the ethics of AI is a topic worth covering in teaching and research.




Figure 3: Topics of particular interest to students (multiple response question, n=422)

4 Assessing the Future Consequences of AI

Narratives and expectations are important factors in the implementation and regulation of new technologies (Sartori & Bocca, 2023). These expectations concern, first and foremost, the benefits and dangers of artificial intelligence for social and personal life. They are crucial components of moral attitudes towards AI. Yet, popular views of the future impact of AI are prone to oversimplifications as either doomsday- or entering-the-golden-age-scenarios. Such dystopian and utopian views contribute to a “myth of AI” (Hoffman et al., 2022), exaggerating the positive or negative impacts, leaving little room for the uncertainty and unpredictability of technological innovation and diffusion (Boyd & Holton, 2018). For these reasons, it is crucial to know how students assess the potential future consequences of AI. Figure 4 presents eight consequences of AI in decreasing order of agreement.

Overall, students have a rather skeptical view of the societal impacts of artificial intelligence. A majority of students agrees with the statement that AI is going to reduce their privacy and is going to take jobs away from people. The strong agreement with the latter is rather astonishing given the fact that the impacts of AI on the labor market and job loss are still heavily debated within academic research and highly uncertain (Arntz et al., 2016; Boyd & Holton, 2018). Students are also rather skeptical or dismissive of the view that AI will improve safety, a widely cited advantage of automation in transportation, medicine, defense, and so on (Bonfanti & Kohler, 2020). At the same time, the proportion of students disagreeing with the statement that AI is going to improve society is twice as large as the proportion of agreement. On average, students clearly do not believe that AI leads to societal progress. However, respondents also see potential benefits of AI. They strongly agree

with the statement that AI is going to make life easier. Thus, from the perspective of the students, AI's positive impact mostly relates to personal convenience.




Figure 4: Future consequences of AI (n=267-282)

A lesser-known issue of AI concerns its environmental impact (Perdue, 2017). As the argument goes, machine learning requires enormous amounts of computational power and data, which is energy intensive and causes environmental emissions (e.g., server cooling). However, the statement that artificial intelligence is going to increase environmental pollution (“Environmental impact”) receives more disapproval than approval. Generally speaking, artificial intelligence is not considered detrimental to ecological sustainability among students.

Finally, with the advancements of artificial intelligence, these technological systems exhibit more and more behaviors previously considered hallmarks of humanity, such as rationality, expert knowledge, or creativity (Stein et al., 2019). What if AI had the same abilities as humans? A clear majority of respondents would find it uncanny if AI had the same capabilities (“Uncanniness”). It is the statement with the strongest agreement overall. However, students are very divided regarding the question whether the importance of human beings would also be reduced. While 38 percent agree, 44 percent disagree, just 18 percent are undecided. While the idea of human-like AI thus causes feelings of uncanniness in most respondents, it does not necessarily pose a threat to their self-esteem as human beings.

5 Moral Evaluation of AI

So far, it has become evident that the accelerated introduction of artificial intelligence raises contentious moral questions. Hence, artificial intelligence is not only judged from a purely instrumental perspective but is also evaluated according to moral categories and standards. That is to say, artificial intelligence becomes morally

embedded, relating to notions of trust, legitimacy, and moral agency (Beckert, 2012; Coeckelbergh et al., 2016; Seo & Lee, 2021).

Our results show that trust is on a low level (Figure 5). Only 26 percent agree that AI developers are trustworthy, while this value drops to 18 percent for AI systems and to 14 percent for companies using artificial intelligence. In fact, more detailed analysis shows that no respondent fully trusts AI systems. Yet, most respondents are unsure whether to trust or not, indicated by the large proportion of neutral responses, especially in the cases of developers and companies. This might reflect the newness of the field and the resulting state of radical uncertainty (Beckert, 2016).

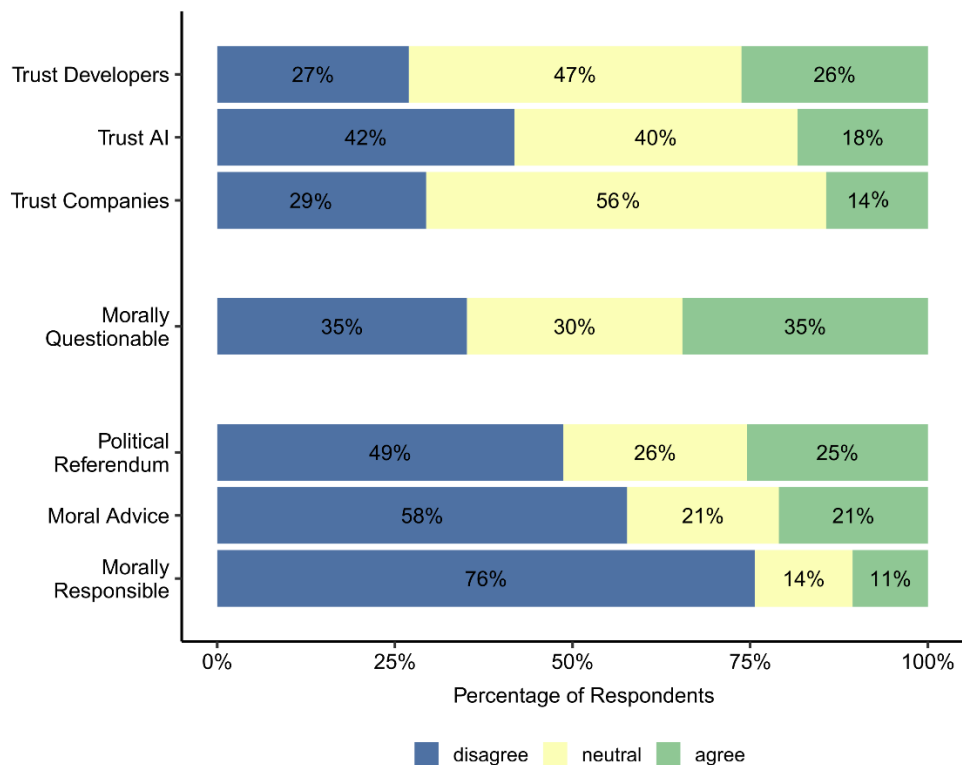


Figure 5: Moral evaluation of AI (n=267-354)

Students are much more divided when it comes to an overall evaluation of the moral legitimacy of artificial intelligence. Just as many respondents find it morally questionable that AI takes on various tasks in society as respondents who don't. In any case, consistent with previous results on the future consequences of AI, students' moral evaluations of AI are not particularly favorable, with substantial variation on this question (Bigman & Gray, 2018).

Artificial intelligence may also have an increasing impact on moral decisions or even select courses of action autonomously, gaining moral agency. In this regard, we find a surprising willingness of students to include AI in moral questions. Around every fifth respondent could very well imagine asking an artificial intelligence for moral advice. Even more striking, every fourth respondent believes that an artificial intelligence would rather be able to recognize which is the better choice in a political referendum than the Swiss population. While agreement is not particularly high in absolute terms, given the unlikelihood and vast implications of these delegations we deem

them especially noteworthy. Moreover, although students are to a certain extent willing to hand over moral and political judgments, they are reluctant to attribute moral responsibility to AI. Just around every 10th respondent agrees that AI could be responsible for its behavior. This peculiar combination of delegation and lack of responsibility raises important questions for moral accountability (Mittelstadt et al., 2016).

6 Moral Rights and Moral Value of AI

Several arguments have been put forward in academic and public discourse for attributing moral value or moral rights to AI. In the legal context, scholars contend copyrights in case of generative AI for text or images (Sundara Rajan, 2019) or even advocate the concept of an electronic person (Matsuzaki & Lindemann, 2016). Philosophers argued that human-likeness (Danaher, 2017), intrinsic moral worth of technological objects (Floridi, 2005), or social relations between AI and human agents (Coeckelbergh, 2010) are sufficient to grant moral rights. Despite these arguments, there is a lack of empirical data on the perceptions of the moral value and rights associated with AI among non-professional academics, including students (Harris & Anthis, 2021).




Figure 6: Attribution of moral value and moral rights (n=263-266)

We asked participants how much they agree with the view that various entities, including AI, have moral value or should possess moral rights. Figure 6 depicts approval in descending order. This allows us to map the positions of different entities within a moral circle, indicating the extension of moral concern to various objects. We find a striking difference between human and non-human animals on the one hand (human and chimpanzees) and non-animate objects (AI, stone, computer program, vending machine) on the other. While the vast majority of respondents attributes value and rights to the former, approval is below 15 percent for all other entities. Looking more closely at the non-animate objects, we find a second, less pronounced distinction. Respondents make a difference between AI and a simple computer program. Approval for the former is more than twice as strong compared to the computer program. It is also clearly stronger than for a common vending machine – a mostly

mechanical device. In the moral circle, AI therefore appears to be closer to humans than a computer program or a vending machine.

These findings certainly lend some credibility to the argument that human-likeness and integration in social relations promote the attribution of moral rights (Coeckelbergh, 2010; Danaher, 2020). Nevertheless, the overall low approval rates indicate that AI still seems to be understood as a mere means to an end. At least from the perspective of students, AI does not possess moral value or rights in a similar way to humans or chimpanzees. They position AI behind those entities in the moral circle. Despite the ongoing discussion by legal scholars and philosophers, a proposal to give artificial intelligence moral rights would only be supported by a very small minority.

7 Conclusion

The continuous arrival of artificial intelligence into more and more areas of daily life confronts social institutions with new challenges (Bostrom, 2014; Mittelstadt et al., 2016). Empirical studies showing the actual use, perceptions, and assessments of AI are essential to effectively tackle these issues. This is also true for universities. A better understanding of the attitudes of university students helps formulating strategies for teaching and regulation (Chan, 2023). It also sheds light on the perspectives of an important social group. Students are key carriers of social activism and innovation in various social areas, from science, to entrepreneurship, to culture and politics (Bernardino & Freitas Santos, 2022; Dahlum & Wig, 2021).

For these reasons, we conducted a full census among students of all faculties at the University of Lucerne in May and June 2023. Nearly all respondents use AI-based technologies, especially intelligent translators or social media. A majority of students have used chatbots (such as ChatGPT) or image recognition software. Yet, despite the widespread use, we observe high levels of uncertainty regarding AI. Students do not feel particularly confident about their knowledge in AI and are undecided whether to trust artificial intelligence or not. The high level of uncertainty is paired with a skeptical perspective of the future consequences of AI, such as violation of privacy, job loss, self-worth, or societal progress. Accordingly, students are particularly interested in moral questions relating to new technologies. Surprisingly, though, nearly every fourth student would be willing to include AI in solving moral and political questions. Only a small minority attributes moral responsibility to AI, raising future questions for moral accountability. Finally, students are reluctant to attribute moral value or moral rights to AI. In sum, we characterize students as:

Highly skeptical users of AI-based technologies, open for its potential in moral and political life but worried about its impact on society.

Below, we point out four major practical recommendations, which can be derived from the results of our research:

1) **Provide information about trustworthy AI:**

A rather surprising finding is the high uncertainty regarding trust in AI. Universities should therefore give information on trustworthy AI and procedures to assess the trustworthiness of AI systems, developers, and providers – in the same way we train students to identify trustworthy sources, publishers, and

literature. In addition, there should be concrete guidance on the legal use of AI applications and on how to correctly indicate AI generated outputs, increasing transparency and scientific integrity.

2) ***Increase literacy regarding AI:***

We observe a high level of insecurity regarding AI-related knowledge among students. This runs counter to the conception of digital natives, assumed to be literate in new technologies. Higher education should therefore overcome such a simplified ideal type. In educational settings, we should not only foster a responsible use of AI but also the ability to understand and explain these new technological resources.

3) ***Create spaces for critical discussion on particular topics:***

Students feel a general need to discuss the morality of AI. Yet, we identified topics deserving special attention. Students should get better acquainted with research on automation and job loss; they should critically reflect on the intersection of ecological sustainability and AI, and they should be able to contemplate the impact of AI on self-worth as human beings.

4) ***Take AI seriously while being realistic:***

Clearly, students use AI-based technologies. But, at least at the time of data collection, popular examples of AI, such as chatbots, were only used regularly by less than one third of the respondents. Other famous applications, such as image generators, represented a small niche. Hence, while AI is in use, we also shouldn't resort to alarmist accounts of the ubiquity of AI. Furthermore, extreme predictions of the future of AI, whether dystopian or utopian, should be addressed with more caution. Taking a closer look at the gray zones of opportunities and challenges not only leads to a more realistic approach towards AI, but also to a more precise assessment of these new technologies.

We also want to point out some limitations to this study. First, we can't rule out self-selection of participants due to the topic of the study ("new technologies – new responsibility?"). Students with a special interest in the morality of new technologies might have been more likely to participate. Second, we are dealing with a technological application that is quickly evolving. Our research provides a snapshot of the situation in summer 2023. It is possible that certain trends, for example the increasing use of various AI-applications, have continued or leveled off. Third, we are only able to make inferences to students at the University of Lucerne and not to students in general, let alone the Swiss population. Last, our primary goal was simply to describe central tendencies and distributions of various aspects related to AI. It is beyond the scope of this paper to analyze correlations or provide theoretical explanations for various outcomes, such as using AI or attributing moral qualities to AI.

In conclusion, this study provides strong evidence for the moral embeddedness of artificial intelligence. As a socio-technological system, the implementation and regulation of AI crucially depends on solid knowledge of the hopes, fears, beliefs, and practices by various social groups. The social sciences, humanities, and related disciplines therefore play an essential role in creating a society where AI can be used for the common good. Students are eager to participate in this discussion.

Appendix A. Data and Methods

To collect data on the use, perceptions, and assessments of AI among students at the University of Lucerne (Switzerland), we conducted a standardized online survey from the 16th of May 2023 to the 14th of June 2023. We sent an E-Mail invitation with a link to all enrolled students. The survey thus represents a full census of the students from all faculties. The title of the survey was “new technologies – new responsibility?”, avoiding any reference to artificial intelligence in order to reduce self-selection of students especially interested in AI. Additionally, we incentivized participation with a prize draw for a cinema voucher. The questionnaire was in German. In total, 440 students participated in the survey, corresponding to a response rate of 12%. Considering the lower response rates of online surveys (Nulty, 2008) and the rather niche topic, we consider this to be acceptable.

Looking at Table 1, providing an overview of the sample composition, we find that female respondents are slightly overrepresented compared to the population of enrolled students at the University of Lucerne (around 60% female students). Culture and social sciences (KSF) are also overrepresented compared to the population (25%), while law and economics are slightly underrepresented (41% and 15% in the population, respectively). On average, respondents are at an advanced stage of their studies and slightly older compared to the general student population. The clear majority lives in Lucerne.

Gender		Age (in years)		Field of study		Semesters		Place of residence	
Female	67%	Min	18	KSF	38%	Min	1	Lucerne	51%
Male	32%	Mean	27	Law	34%	Mean	7	Zurich	10%
Other	1%	Max	77	Economics	10%	Max	21	Aargau	5%
				Theology	10%			Bern	4%
				Health sciences	8%			Other	30%
n	345	n	331	n	341	n	336	n	314

Table 1: Description of the sample (KSF=Humanities and Social Sciences)

At the beginning of the questionnaire, participants received instructions and were asked for informed consent. After some introductory questions, we provided a short definition of the term “artificial intelligence”. The next part of the survey was divided into four thematic blocks, covering (i) use of and attitudes towards technology, (ii) ontological categorization of AI, (iii) meta-ethical orientations, and (iv) personal values. Additionally, we asked respondents questions on the moral evaluation of AI, such as moral acceptance, AI rights, or AI trust, in a separate block. To reduce sequential effects, we randomized the order of these blocks and the order of items within these blocks. We measured socio-demographic information at the end.

For the four thematic blocks (i) to (iv) described above, we used a missing-by-design approach (Enders, 2010). Each respondent received three out of the four blocks at random. This produces 25% missing values for these questions by design, resulting in a lower number of available responses (around 100 respondents less). However, it has the considerable advantage of decreasing the length of the survey and avoiding fatigue. The resulting average response time was around 16 minutes. The survey was carried out in Qualtrics.

For the analysis, we simplified the response categories of variables with five-point Likert scales (this concerns Figures 2, 4, 5, and 6). The answers from the extreme poles were aggregated, resulting in three response categories. We combined “strongly disagree” and “somewhat disagree” to “disagree”. Likewise, we combined

“strongly agree” and “somewhat agree” to “agree”. All calculations were performed using R 4.3.0. The exact wording of the questions and the R-code to reproduce the calculations and graphs are available on request.

References

- Almaiah, M. A., Alfaisal, R., Salloum, S. A., Hajje, F., Thabit, S., El-Qirem, F. A., Lutfi, A., Alrawad, M., Al Mulhem, A., Alkhdour, T., Awad, A. B., & Al-Marouf, R. S. (2022). Examining the Impact of Artificial Intelligence and Social and Computer Anxiety in E-Learning Settings: Students' Perceptions at the University Level. *Electronics*, 11(22), Article 22. <https://doi.org/10.3390/electronics11223662>
- Arntz, M., Gregory, T., & Zierahn, U. (2016). The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis: OECD Social, Employment and Migration Working Papers, No. 189. OECD Publishing. <https://doi.org/10.1787/5jlz9h56dvq7-en>
- Beckert, J. (2012). Die sittliche Einbettung der Wirtschaft. Von der Effizienz- und Differenzierungstheorie zu einer Theorie wirtschaftlicher Felder. *Berliner Journal Für Soziologie*, 22(2), 247–266. <https://doi.org/10.1007/s11609-012-0187-y>
- Beckert, J. (2016). *Imagined futures: Fictional expectations and capitalist dynamics* / Jens Beckert. Harvard University Press.
- Bernardino, S., & Freitas Santos, J. (2022). Cultural Differences in Crowdfunding: A Four-Country Study. *E-Revista De Estudos Interculturais Do CEI-ISCAP*, 10.
- Bigman, Y. E., & Gray, K. (2018). People are averse to machines making moral decisions. *Cognition*, 181, 21–34. <https://doi.org/10.1016/j.cognition.2018.08.003>
- Bonfanti, M. E., & Kohler, K. (2020). Artificial Intelligence for Cybersecurity. <https://doi.org/10.3929/ethz-b-000417116>
- Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies* (First edition). Oxford University Press. <http://site.ebrary.com/lib/alltitles/docDetail.action?docID=10896241>
- Boyd, R., & Holton, R. J. (2018). Technology, innovation, employment and power: Does robotics and artificial intelligence really mean social transformation? *Journal of Sociology*, 54(3), 331–345. <https://doi.org/10.1177/1440783317726591>
- Chan, C. K. Y. (2023). A comprehensive AI policy education framework for university teaching and learning. *International Journal of Educational Technology in Higher Education*, 20(1), 38. <https://doi.org/10.1186/s41239-023-00408-3>
- Coeckelbergh, M. (2010). Robot rights? Towards a social-relational justification of moral consideration. *Ethics and Information Technology*, 12(3), 209–221. <https://doi.org/10.1007/s10676-010-9235-5>
- Coeckelbergh, M., Pop, C., Simut, R., Peca, A., Pintea, S., David, D., & Vanderborght, B. (2016). A Survey of Expectations About the Role of Robots in Robot-Assisted Therapy for Children with ASD: Ethical Acceptability, Trust, Sociability, Appearance, and Attachment. *Science and Engineering Ethics*, 22(1), 47–65. <https://doi.org/10.1007/s11948-015-9649-x>
- Dahlum, S., & Wig, T. (2021). Chaos on Campus: Universities and Mass Political Protest. *Comparative Political Studies*, 54(1), 3–32. <https://doi.org/10.1177/0010414020919902>

- Danaher, J. (2017). The Symbolic-Consequences Argument in the Sex Robot Debate. In J. Danaher & N. McArthur (Hrsg.), *Robot Sex: Social and Ethical Implications*. MIT Press. <https://philarchive.org/rec/DANTSA-6>
- Danaher, J. (2020). Welcoming Robots into the Moral Circle: A Defence of Ethical Behaviourism. *Science and Engineering Ethics*, 26(4), 2023–2049. <https://doi.org/10.1007/s11948-019-00119-x>
- Enders, C. K. (2010). Applied missing data analysis. *Methodology in the social sciences*. Guilford.
- Floridi, L. (2005). Information ethics, its nature and scope. *ACM SIGCAS Computers and Society*, 35(2), 3. <https://doi.org/10.1145/1111646.1111649>
- Gherheș, V. (2018). Artificial Intelligence: Perception, expectations, hopes and benefits. *Romanian Journal of Human - Computer Interaction*, 11(3), 220–231.
- Harris, J., & Anthis, J. R. (2021). The Moral Consideration of Artificial Entities: A Literature Review. *Science and Engineering Ethics*, 27(4), 53. <https://doi.org/10.1007/s11948-021-00331-8>
- Hoffman, S. G., Joyce, K., Alegria, S., Bell, S. E., Cruz, T. M., Noble, S. U., Shestakofsky, B., & Smith-Doerr, L. (2022). Five Big Ideas About AI. *Contexts*, 21(3), 8–15. <https://doi.org/10.1177/15365042221114975>
- Litvinski, O. (2023). On Algorithmic Opacity and Moral Certainty. <https://doi.org/10.31235/osf.io/rjsvk>
- Matsuzaki, H., & Lindemann, G. (2016). The autonomy-safety-paradox of service robotics in Europe and Japan: a comparative analysis. *AI & SOCIETY*, 31(4), 501–517. <https://doi.org/10.1007/s00146-015-0630-7>
- Mihálka, R. (2023). How can universities use AI? And what do they have to keep in mind while using AI? November 27. Event at the University of Lucerne, Switzerland.
- Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 205395171667967. <https://doi.org/10.1177/2053951716679679>
- Nulty, D. D. (2008). The adequacy of response rates to online and paper surveys: What can be done? *Assessment & Evaluation in Higher Education*, 33(3), 301–314. <https://doi.org/10.1080/02602930701293231>
- Perdue, R. T. (2017). Superintelligence and Natural Resources: Morality and Technology in a Brave New World. *Society & Natural Resources*, 30(8), 1026–1031. <https://doi.org/10.1080/08941920.2016.1264652>
- Prensky, M. (2001). Digital Natives, Digital Immigrants. *On the Horizon*, 9(5), 1–6.
- Sartori, L., & Bocca, G. (2023). Minding the gap(s): public perceptions of AI and socio-technical imaginaries. *AI & SOCIETY*, 38(2), 443–458. <https://doi.org/10.1007/s00146-022-01422-1>
- Seo, K. H., & Lee, J. H. (2021). The Emergence of Service Robots at Restaurants: Integrating Trust, Perceived Risk, and Satisfaction. *Sustainability*, 13(8), 4431. <https://doi.org/10.3390/su13084431>
- Stein, J.-P., Liebold, B., & Ohler, P. (2019). Stay back, clever thing! Linking situational control and human uniqueness concerns to the aversion against autonomous technology. *Computers in Human Behavior*, 95, 73–82. <https://doi.org/10.1016/j.chb.2019.01.021>
- study.com. (2023). Productive Teaching Tool or Innovative Cheating? Study.Com. <https://study.com/resources/perceptions-of-chatgpt-in-schools>
- Sundara Rajan, M. T. (2019). Moral rights: The future of copyright law? *Journal of Intellectual Property Law & Practice*, 14(4), 257–258. <https://doi.org/10.1093/jiplp/jpz008>