

Mind the Gap: Positioning Generative AI in Higher Education – A Comparative Study of Academic Staff and Students' Attitudes

Wouter de Raad , Ewa Stefania Malinowska , Karolina Małek , Agnieszka Wojnarowska 

Department of Psychology, University of Warsaw, Poland.

How to cite: de Raad, W., Malinowska, E. S., Małek, K., Wojnarowska, A. (2026). Mind the Gap: Positioning Generative AI in Higher Education – A Comparative Study of Academic Staff and Students' Attitudes. In: 12th International Conference on Higher Education Advances (HEAd'26). Valencia, 15-18 June 2026. <https://doi.org/10.4995/HEAd26.2026.21701>

Abstract

This study investigates the adoption of Generative AI at the University of Warsaw, a major European public university. Integrating Large Language Models (LLMs) into higher education offers pedagogical potential but threatens academic integrity. Despite AI's ubiquity, institutional responses remain fragmented and inconsistent. This comparative study examines how usage patterns, attitudes, and perceived risks of Generative AI differ between students and academic staff within the same university.

Findings reveal a counterintuitive image where faculty adoption rates exceed those of students. A significant communication "dead zone" exists, however: students clearly perceive restrictive policies but fail to recognize faculty encouragement of AI. Knowledge transfer is horizontal (peer-to-peer) rather than vertical. Whereas both groups view AI as a creative scaffold, only faculty struggle with feeling overwhelmed by the sudden and rapid advance of AI and a self-imposed responsibility to prevent student "cognitive atrophy." The university operates under a "hidden curriculum" where AI is utilized but pedagogically marginalized. Policies must shift from prohibitions toward frameworks integrating AI literacy into formal mentor-student relationships.

Keywords: Generative AI, Higher Education, Student-Teacher Approach Comparison, AI Adoption

1. Introduction

Artificial intelligence has been increasingly recognized in recent years as a technology with the potential to fundamentally transform higher education (Bates et al., 2020; Zhang & Aslan, 2021). However, the sudden emergence and widespread availability of Large Language Models (LLMs), such as ChatGPT, in late 2022, led to transformative changes in academic education. Unlike previous educational technologies that were implemented top-down by institutions, AI

adoption occurred bottom-up (Valentini & Blancas, 2025)—driven primarily by students seeking efficiency. Currently, universities worldwide face the urgent challenge of redefining teaching, learning, and assessment processes in a reality where academic-level content generation is available at one's fingertips (Sullivan et al., 2023). The rapid development of AI tools justifies the need to examine evolving attitudes, usage patterns, and perceived benefits and risks as a foundation for developing evidence-based regulatory frameworks (Barros et al., 2023; Chan, 2023). As indicated by recent UNESCO reports (Sabzalieva & Valentini, 2023; UNESCO, 2024), universities must rapidly transition from a stage of "shock" to systemic adaptation to fully realize the potential of AI in higher education. However, institutional reactions are often fragmented, focusing primarily on usage guidelines and plagiarism issues rather than comprehensive skill development (Valentini & Blancas, 2025).

2. Literature review

The introduction of GenAI into the academic ecosystem carries dual implications. On one hand, the literature points to numerous opportunities: learning personalization, support in idea generation, and the automation of routine tasks (Michel-Villarreal et al., 2023). On the other hand, researchers raise alarming concerns regarding academic integrity, model hallucinations, and the risk of eroding critical thinking skills (Farrelly & Baker, 2023). The key issue becomes how different user groups—students and teachers—perceive this cost-benefit balance.

2.1. The student perspective: pragmatism and risk

Compared to institutions, students typically adapt to new technologies more rapidly, utilizing LLMs as "digital assistants" for brainstorming, summarizing complex texts, or improving writing style to boost productivity (Chan & Hu, 2023). However, they remain cautious regarding reliability (model hallucinations) and academic ethics (Ravselj et al., 2025). At the same time, students highlight the limitations of AI tools in areas such as critical thinking, numeracy skills, and interpersonal communication. In the context of higher education, this creates the risk that, without proper guidance, students may become over-dependent on technology, potentially missing opportunities to independently engage in complex cognitive activities.

2.2. The faculty perspective: skepticism, need for support, and the institutional gap

Conversely, academic teachers often approach innovation with greater reserve. Faculty members face ethical challenges and the necessity of verifying learning outcomes in conditions where the authorship of work becomes unclear (Farrelly & Baker, 2023). To effectively integrate AI, instructors require trust and institutional support, which is currently lacking (Chounta et al., 2022). The absence of clear guidelines and training leaves many instructors feeling unsure in this new reality, which may lead to defensive attitudes or a complete disregard

for AI in the didactic process. This divergence between widespread student adoption and limited instructor engagement creates a significant gap requiring immediate policy intervention (Michel-Villarreal et al., 2023).

3. Study objective: identifying the gap

The present study aims to present this gap through a comparative analysis of two key stakeholder groups within a single academic ecosystem—the University of Warsaw. The study provides insight into how the perspectives of students and staff differ and on what points they converge. We compare the attitudes, usage patterns, and perceptions of GenAI among students (Faculty of Psychology, UW) and academic teachers at the University of Warsaw. By juxtaposing the perspectives of these two groups, we aim to identify the gap in expectations and practices, which constitutes a necessary step for developing consistent and effective educational and regulatory recommendations.

Based on the identified gaps, this study addresses the following research question:

RQ1: To what extent do usage patterns, attitudes, and perceived risks of Generative AI differ between students and faculty members?

4. Methods

4.1. Participants and Procedure

Two surveys were conducted. The first study was carried out at the end of 2024 among students at the Faculty of Psychology, University of Warsaw. Two surveys were conducted, starting with a 2024 study of psychology students at the University of Warsaw ($N=113$; $Mage=21.59$, $SD=2.97$). The sample was predominantly female ($n=80$), with 26 men and 7 identifying as "Other" or undisclosed. Participants represented all years of a five-year Master's program, with a balanced distribution across the first three years (approx. 24% each) and consistent representation in the senior years (13% each). This diverse cohort captures perspectives ranging from early-stage students to those nearing professional qualification. The second study, conducted in spring 2025, involved 178 University of Warsaw teaching staff ($Mage=43.55$, $SD=11.25$, range 26–71). The cohort comprised 95 women, 73 men, and 10 who identified as "Other" or undisclosed. Disciplines were led by Social Sciences ($n=54$), followed by Humanities ($n=24$), and Natural/Exact Sciences ($n=22$), with additional representation from Linguistics, Engineering, and Arts. Professionally, the sample was dominated by assistant professors ($n=103$), alongside full professors ($n=32$), PhD students ($n=20$), assistants ($n=9$), and 14 external or undisclosed instructors. Participants were external university teachers or did not reveal their professional position. Both groups were invited to fill out an anonymous online

survey, sent via internal email. It contained open questions about the methods of AI utilization, their assessment of its utility, reasons for non-adoption, and the patterns of communication concerning AI within the learning/teaching context of the academy.

5. Analysis

In this study, comparative quantitative analysis and exploratory content analysis procedures were employed to investigate how students and academic staff conceptualize the presence and role of artificial intelligence within the academic environment. The analysis focused on identifying how these groups position AI in the context of learning, teaching, and professional work, based on open-ended questions regarding the use of AI in various academic scenarios, such as preparing for classes, studying for exams, and completing written assignments or projects. We also investigated the ways AI usage is being communicated between the groups. For the content analysis reliability, three independent researchers acted as competent judges. These researchers performed an inductive analysis. The coding focused on two primary dimensions: the functional application of AI tools (e.g., explaining complex theories, automating technical tasks, or generating creative inspiration) and the axiological positioning of AI (e.g., as a "useful but fallible assistant" or a "threat to academic integrity").

6. Results

6.1. AI adoption rates: Faculty vs Students

Contrary to the common assumption that students are the primary users of new technologies, the data reveal that academic teachers reported higher usage rates than students.

Table 1. AI usage rates in the last semester

Group	Used AI (%)
Faculty	74.2%
Students	65.5%

6.2. Communication gap

There is a stark difference in how often staff report that they discuss AI in the classroom with students, and the frequency with which students report that this happens. While faculty members believe they are addressing the topic, students do not seem to perceive it as frequently. 37.6%

of lecturers claim they "Often" referred to AI use by students during the semester, whereas only 19.5% of students felt their lecturers referred to AI "Often".

The study further highlights a mismatch in perceived "encouragement" to use these tools. 42.1% of lecturers claim they encouraged students to use AI, but only 20.4% of students actually felt encouraged. The data on bans is more closely aligned; 45.5% of faculty reported banning AI in certain areas, while 49.6% of students reported encountering such bans. This suggests that while prohibitions are communicated clearly, supportive or constructive guidance is often lost in translation. The transfer of AI knowledge within the university appears to be "dead," with both groups relying heavily on external sources or peer - to - peer learning rather than each other. The university does not educate students when it comes to AI.

6.3. The failure of vertical transfer

The transfer of AI knowledge within the university hierarchy appears to be minimal. As detailed in Table 2, both groups rely heavily on external sources or horizontal peer-to-peer learning rather than on academic exchange.

Table 2. Primary sources of information regarding AI tools

Source of knowledge	Students (%)	Faculty (%)
Independent Internet Search	63.7%	88.8%
Peer Group (Students/Other Lecturers)	51.3%	46.1%
Social Media	50.4%	20.2%
Friends outside the University	49.6%	45.5%
Vertical Academic Channel (Lecturers <-> Students)	6.2%	12.4%

A comparative analysis of knowledge sources demonstrates a strong reliance on horizontal, peer-to-peer expertise within each demographic. Faculty members appear to lean primarily on internal professional networks, with 46.1% reporting that they acquire AI knowledge from other academics within the university. Conversely, students demonstrate a similar reliance on their

own social circles, with 51.3% citing other students and 49.6% citing friends outside the university as primary sources of information.

However, the vertical communication channel—the traditional mentor-student relationship — is remarkably absent in the context of AI. Even though 74.2% of faculty and 65.5% of students are active AI users, only 6.2% of students report learning about these tools from their lecturers. Similarly, only 12.4% of faculty members acquire AI insights from their students. This data suggests that while both groups are independently building their AI literacy through internet research and peer interaction, the formal university hierarchy remains a "dead zone" for knowledge transfer, leaving the two cohorts to navigate the AI landscape in isolation from each other.

6.4. Qualitative results: patterns of use and perception

The following section presents the summary of results of a content analysis of statements provided by students (Study 1) and academic members (Study 2). In response to open-ended questions, we sought to determine how students and academic staff perceive AI tools: as assistants, educators, intruders, or disruptors? Due to space constraints and the extensive nature of the qualitative findings, only a summary of the results is presented here.

The first major alignment between group perspectives appears in Parallel 1: The Private Tutor (Students) and the Object of Critical Analysis (Teachers). Both groups recognize AI's potential to facilitate deep learning. Students adopt the tool as a "learning sparring partner" to explain complex jargon, while faculty strategically mirror this by using AI's inaccuracies as didactic tools to foster evaluative thinking. This suggests that whereas students use AI to bridge gaps in understanding, faculty utilize those very gaps to sharpen critical literacy. This utility extends to Parallel 2: The Rigorous Editor and Technical Support, where AI is framed as a "functional optimizer." Both cohorts relegate "labor-intensive, uncreative" tasks to the technology—students focus on syntax and structural coherence, while faculty automate "behind-the-scenes" logistics like syllabi creation and administrative correspondence. Similarly, Parallel 3: The Creative Spark and Content Architect highlights AI's role as "conceptual scaffolding." Whether overcoming writer's block or generating raw didactic materials, both groups treat AI as a foundational springboard rather than a source of finished intellectual products. A significant finding is the shared "limited trust" model. Both groups exhibit high awareness of AI's "credibility deficit," specifically regarding confabulation and the fabrication of sources. Students characterize AI as a "useful but unintelligent assistant," a sentiment echoed by university teachers who view it as a "parasitic entity" that lacks logical depth. This shared skepticism forms the basis of Parallel 4, where motivations for distancing—such as the perception of AI as a "noise generator" or an "incompetent assistant"—align around the pragmatics of time-efficiency and substantive quality.

Despite these parallels, the faculty perspective diverges in two critical dimensions. Unlike students, who primarily cite pragmatic or ethical reasons for non-adoption, faculty express a distinct sense of being overwhelmed. AI is viewed as an "Inaccessible Novelty"; instructors report a lack of "bandwidth" and specialized training, viewing the technology as a daunting addition to an already high academic workload.

Additionally, though both groups identify ethical risks, only academic staff perceive this as a systemic threat. They perceive AI as an "ethical usurper" that threatens the foundations of critical thinking. Consequently, faculty increasingly view themselves as the guardians of knowledge, responsible for maintaining the integrity of academia.

7. Discussion

The findings of this study reveal an intricate, dynamic, and often counterintuitive landscape of AI adoption at the University of Warsaw. Contrary to the "technophobic" teacher and the "digital native" student, it is the academic staff that reports higher usage rates than students, yet they simultaneously perceive AI as a formidable barrier requiring specialized, complex training. They perceive the necessity of upholding the integrity of the "learning process" against the "greed of big tech", and the potential "intellectual atrophy" of their students, yet a profound communicative and pedagogical gap prevents the integration of these perspectives into a shared academic culture.

7.1. Restrictions vs Encouragement

One of the most striking findings is the asymmetry in communication effectiveness. The data shows a consensus regarding prohibitions: roughly half of both groups acknowledge the existence of AI bans. However, a large discrepancy exists regarding encouragement. Whereas over 40% of teachers claim to encourage AI use, only 20% of students perceive this. Maybe restrictive communication is clearly perceived, but encouraging words are not? Prohibitions are binary, clear, and easy to understand and transmit. Encouragement, on the other hand, without clear guidelines or training, could be perceived by students as vague or non-existent.

This establishes a hidden curriculum in which artificial intelligence is viewed primarily through the lens of academic integrity and risk, rather than as a legitimate tool for fostering intellectual development. The qualitative analysis corroborates this finding: students frequently perceive AI as a private tutor or a covert editor used in isolation, rather than a subject of transparent, open discussion within the classroom setting, even though faculty members themselves report efforts to integrate these tools into their teaching practice.

7.2. The “dead zone” of vertical knowledge transfer

At the time of this study, the University of Warsaw lacked a unified, top-down AI policy, leaving regulations to the discretion of individual faculties or instructors. This decentralized approach contextualizes our findings. Perhaps the most alarming insight from this study is the identification of a knowledge transfer dead zone within the institutional hierarchy. Despite high adoption rates in both groups, the vertical transmission of expertise—from mentor to student—is virtually non-existent. Instead, the university functions within two parallel spheres of horizontal learning: students navigate the ethical and practical complexities of AI through peer interaction and internet research, while faculty develop their own competencies in isolation. This disconnection exacerbates the digital dissonance, as neither group acts as a resource for the other, leaving the formal academic structure largely irrelevant in the process of building AI literacy.

7.3. Resistance and distancing: a potential for shared ethical dialogue

The alignment between students and faculty in distancing themselves from AI stems from a shared view of the tool as both a pragmatic “noise generator” and an ethical intruder. However, a key divergence emerges in their professional positioning: faculty frame themselves as guardians of the academic ethos. Perhaps this heavy sense of responsibility likely drives their perception of AI as a daunting, “inaccessible novelty” requiring extensive formal training, whereas students—though not careless—approach the technology with greater functional flexibility. The shared capacity for meta-reflection on “intellectual atrophy” and academic integrity provides a robust ethical foundation for a unified institutional stance on AI. However, this potential can only be realized if the “hidden curriculum” is replaced by open dialogue, bridging the gap between faculty’s ethical guardianship and students’ pragmatic exploration to foster a culture of shared responsibility.

8. Conclusion and recommendations

Encouragement must be replaced with specific assignments where AI use is mandatory but critically analyzed (e.g., find the hallucination).

Vague encouragement must be replaced by credit-bearing assignments where AI use is mandatory but serves as a primary object of critical deconstruction. Through structured tasks like “hallucination audits,” students move from passive users to analytical evaluators, directly achieving learning outcomes related to information literacy and cognitive independence. In this framework, the teacher’s role shifts from a mere gatekeeper to a proactive facilitator who actively encourages AI experimentation. Guiding students through deliberate exercises, the instructor ensures they perceive AI not as a shortcut, but as a catalyst for developing the critical skepticism and ethical rigor essential for professional practice.

Since peer learning is dominant, universities should create mixed student-faculty working groups to share AI strategies openly.

Second, to overcome the failure of vertical knowledge transfer, the university must leverage the strength of existing horizontal networks. Since both cohorts rely heavily on peer-to-peer learning, the establishment of mixed student-faculty working groups is recommended. These collaborative forums would allow for the open sharing of AI strategies, de-stigmatizing the technology and aligning the faculty's role as "guardians of knowledge" with the students' practical needs. By fostering a culture of mutual experimentation, the academy can mitigate faculty's competence-related anxieties and ensure that AI serves as a catalyst for cognitive independence rather than a substitute for it.

Academic teachers' support.

To address the prevailing sense of being overwhelmed by "inaccessible novelty," universities must shift their focus from purely technical training toward student-centered learning and workshops centered on pedagogical strategies. Rather than expecting every faculty member to achieve AI technical mastery, institutions should provide ready-to-use didactic frameworks—such as lesson plans where AI serves as a "devil's advocate" to stimulate critical thinking. By offering these practical, classroom-tested examples, institutions can alleviate professional anxiety and reinforce the educator's essential role as a critical guide in the digital era.

9. Limitations

Several limitations should be considered. First, the sample structure presents an asymmetry: the student cohort is drawn solely from the Faculty of Psychology, while the faculty sample represents a broader university cross-section. This disciplinary mismatch may influence comparisons regarding ethical attitudes. Second, the voluntary survey format introduces self-selection bias, potentially overrepresenting tech-savvy individuals and inflating adoption rates. Third, reliance on self-reported data risks social desirability bias and the illusion that faculty overestimate the clarity of their instructions. Additionally, semantic ambiguity regarding *using AI* (administrative support vs. content generation) complicates direct usage comparisons. Finally, as a cross-sectional study, these findings represent a specific moment in time within a rapidly evolving technological landscape.

References

- Barros, A., Prasad, A., & Śliwa, M. (2023). Generative artificial intelligence and academia: Implications for research, teaching, and service. *Management Learning*, 54(5), 597–604. <https://doi.org/10.1177/13505076231201445>

- Bates, T., Cobo, C., Mariño, O., & Wheeler, S. (2020). Can artificial intelligence transform higher education? *International Journal of Educational Technology in Higher Education*, 17, 42. <https://doi.org/10.1186/s41239-020-00218-x>
- 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>
- Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *International Journal of Educational Technology in Higher Education*, 20(1), 43. <https://doi.org/10.1186/s41239-023-00411-8>
- Chounta, I.-A., Bardone, E., Raudsep, A., & Pedaste, M. (2022). Exploring Teachers' Perceptions of Artificial Intelligence as a Tool to Support their Practice in Estonian K-12 Education. *International Journal of Artificial Intelligence in Education*, 32, 725–755. <https://doi.org/10.1007/s40593-021-00243-5>
- Farrelly, T., & Baker, N. (2023). Generative artificial intelligence: Implications and considerations for higher education practice. *Education Sciences*, 13(11), 1109. <https://doi.org/10.3390/educsci13111109>
- Michel-Villarreal, R., Vilalta-Perdomo, E., Salinas-Navarro, D. E., Thierry-Aguilera, R., & Gerardou, F. S. (2023). Challenges and opportunities of generative AI for higher education, as explained by ChatGPT. *Education Sciences*, 13(9), 856. <https://doi.org/10.3390/educsci13090856>
- Ravšelj D, Keržič D, Tomažević N, Umek L, Brezovar N, A. Iahad N, et al. (2025). Higher education students' perceptions of ChatGPT: A global study of early reactions. *PLoS ONE* 20(2): e0315011. <https://doi.org/10.1371/journal.pone.0315011>
- Sabzalieva, E., & Valentini, A. (2023). *ChatGPT and artificial intelligence in higher education: Quick start guide*. UNESCO. <https://unesdoc.unesco.org/ark:/48223/pf0000385146>
- Sullivan, M., Kelly, A., & McLaughlan, P. (2023). ChatGPT in higher education: Considerations for academic integrity and student learning. *Journal of Applied Learning and Teaching*, 6(1). <https://doi.org/10.37074/jalt.2023.6.1.17>
- UNESCO. (2024). *Guidance for generative AI in education and research*. Paris: UNESCO.
- Valentini, A., & Blancas, A. (2025). *The challenges of AI in higher education and institutional responses: Is there room for competency frameworks?* UNESCO International Institute for Higher Education in Latin America and the Caribbean (IESALC).
- Zhang, K., & Aslan, A. B. (2021). AI technologies for education: Recent research & future directions. *Computers and Education: Artificial Intelligence*, 2, 100025. <https://doi.org/10.1016/j.caeai.2021.100025>