

**Stefanija Sikijovska**  
University of Primorska,  
Faculty of Management  
68231028@student.upr.si

**Sofia Vitoria Diks Purdom**  
University of Primorska,  
Faculty of Management  
68230155@student.upr.si

**Iva Tancheska**  
University of Primorska,  
Faculty of Management  
68231063@student.upr.si

**Anastasija Naumcheska**  
University of Primorska,  
Faculty of Management  
68231042@student.upr.si

**Antonio Naumcheski**  
University of Primorska,  
Faculty of Management  
68231004@student.upr.si

**Monika Gichevska**  
University of Primorska,  
Faculty of Management  
68231033@student.upr.si

**Isidora Stefanoska**  
University of Primorska,  
Faculty of Management  
68231076@student.upr.si

**Nikita Vajdeska**  
University of Primorska,  
Faculty of Management  
68231078@student.upr.si

**Violeta Grcheva**  
University of Primorska,  
Faculty of Management  
68231068@student.upr.si

**Igor Rižnar**  
University of Primorska,  
Faculty of Management  
igor.riznar@upr.si

# Exploring the Landscape of AI and LLMs: Familiarity, Usage, and Education Gaps at a Slovenian Hei

This paper aims at gaining insight about the familiarity with AI and LLMs, their usage, and education gaps of students at a Slovenian higher education institution (University of Primorska) and, based on the analysis of the questionnaire results, proposing the necessary future endeavours to raise awareness of pros and cons of using tools like ChatGPT in an educational environment. The paper is organized as follows: after a brief introduction, we give a concise literature review, followed by the research methodology, results of a questionnaire analysis and a conclusion, together with future research recommendations.

*Keywords:* AI ethics, Artificial Intelligence (AI), AI literacy, bias in AI systems, ChatGPT

## **OJS: poznavanje, uporaba in izobraževalne vrzeli na slovenski visokošolski ustanovi**

V članku skušamo pridobiti vpogled v poznavanje umetne inteligence (UI) in velikih jezikovnih modelov (angl. *large language models* – LLM), njihovo uporabo ter vrzeli v izobraževanju študentov na slovenskem visokošolskem zavodu (Univerza na Primorskem) in na podlagi analize rezultatov vprašalnika predlagati potrebna prihodnja prizadevanja za ozaveščanje o prednostih in slabostih uporabe orodij, kot je ChatGPT, v izobraževalnem okolju. Članek je organiziran na naslednji način: po kratkem uvodu podajamo strnjen pregled literature na obravnavanem področju, ki mu sledi predstavitev metodologije, ki smo jo uporabili v raziskavi, zatem podamo rezultate analize odgovorov, postavljenih v vprašalniku. Zapis strnemo z zaključkom ter s priporočili za prihodnje raziskave.

*Ključne besede:* UI-etika, UI-pismenost, umetna inteligenca, predsodki sistemov UI, ChatGPT

 <https://doi.org/10.26493/1854-4231.20.23-29>

Note: This article is the result of a research-based teaching and learning project conducted during the 2023–24 academic year with a group of first-year undergraduate students at the Faculty of Management, University of Primorska. Any mistakes or inaccuracies in the text are solely the responsibility of the teacher overseeing the project. Despite occasional challenges along the way, the experience was an extremely rewarding one, both for the students and the instructor. It offered valuable insights into the process of collaborative research and academic writing, highlighting the importance of critical thinking, feedback, and iterative revision. Through this project, students not only deepened their understanding of the subject matter but also developed some essential academic skills that will benefit them throughout their studies.

## Introduction

The rapid development of artificial intelligence (AI) has generated significant interest in exploring its potential applications and impacts across different fields. The most notable of these advancements has been the introduction of large language models. One of the most famous examples is ChatGPT which was made widely accessible by OpenAI in November 2022. ChatGPT quickly gained popularity, reaching five million users within five days of its launch.

At the core of ChatGPT and similar models lies the concept of generative AI, which is designed to create coherent and contextually relevant text based on the input it receives. LLMs, like ChatGPT, are built on the architecture of Generative Pre-Trained Transformers (GPTs), which leverage vast amounts of data to predict and generate text by identifying patterns and relationships within the language. This capability is facilitated by the transformer model, introduced in 2017, which allows for more efficient processing of context compared to previous models. Despite their sophistication, it is important to note that LLMs do not learn or think like humans. Their “understanding” of language is statistical rather than cognitive.

In education, AI-based applications like ChatGPT have been proposed as tools to enhance learning and teaching by assisting in information searches, generating multilingual content, preparing teaching materials, and aiding in data analysis. Proponents argue that these tools can revolutionize traditional educational practices, offering personalized learning experiences and fostering essential skills. However, concerns have been raised regarding issues such as plagiarism, cheating, and the potential disruption of natural learning processes. These disruptions have the ability to stunt students’ writing skills and their critical thinking abilities.

This paper focuses on understanding the familiarity and usage associated with AI and LLMs (mostly ChatGPT) within the context of a higher education institution in Slovenia (University of Primorska). By conducting a survey among 90 students, our study aims to shed light on students’ perceptions and experiences with ChatGPT, highlighting both the opportunities and challenges posed by its integration into academic settings.

## Literature Review

In *Deep Learning*, John D. Kelleher (2019) provides a concise and accessible introduction to deep

learning, focusing on how it has revolutionized artificial intelligence. He explains the fundamental concepts – like neural networks, backpropagation, convolutional networks, and recurrent networks – in a way that balances technical accuracy with readability. Kelleher also situates deep learning within the broader history of AI, highlighting why and how it became so dominant in the 2010s. Kelleher’s major contribution is in demystifying deep learning for readers who are not AI specialists. He emphasizes the mechanisms by which deep learning systems learn from large datasets and why they outperform traditional machine learning techniques in many fields such as image recognition, natural language processing, and game playing. Importantly, he frames deep learning not as a complete solution to AI, but as a powerful tool with specific strengths and limitations. Kelleher’s introductory book also touches on ethical concerns without deep exploration, such as biases embedded in training data or environmental costs of training large models.

The applicability of LLMs extends across various domains, including customer operations, marketing, sales, software engineering, and research and development. While it demonstrates satisfactory performance in qualitative data tasks, its efficacy in quantitative data analysis and methodology writing is contingent on accurate input information (Mahuli et al. 2023).

While it is apparent that ChatGPT can offer a wide range of assistance, its reliability poses a significant question that requires careful consideration. Even though it can help with qualitative analysis if you give it written information, it has trouble putting everything together in a way that makes sense, and may not give accurate citations.

Critical thinking is central to writing and producing an accurate academic paper and it is a great risk that ChatGPT is incompetent in this area. A greater problem is created when students rely on AI to summarize, write, and understand text from a model that cannot think critically. Because of dependence on ChatGPT, more and more people do not know how to write, summarize, or even paraphrase. When we have AI to do the work for us, we lose motivation to do it ourselves. Additionally, AI is used to create content in areas like journalism, marketing, academia, and even in law. The concern is academic honesty. If someone uses AI to write a paper without giving credit, it can lead to serious problems. This way, people can share misleading information because we do not know if ChatGPT gives us accurate

information. This is because ChatGPT reads a lot of text before answering a question and not all sources are reliable. So, even though AI tools are very useful, it is important to know when something is written by AI and when it is not (Rahman et al. 2023).

Despite occasional fluency, ChatGPT lacks human-like understanding and may produce inaccurate responses, which facilitates the spread of misinformation (Farhat et al. 2023).

The literature on artificial intelligence (AI), AI assistants, and ChatGPT reveals a dynamic and multifaceted exploration of technology's capabilities, limitations, and implications for society. Scholars, technologists, and ethicists converge on several critical themes: the specialized versus general nature of AI intelligence, the societal risks of overreliance, and the ethical frameworks necessary for its responsible use.

Igor Rižnar (2024) examines ChatGPT from a linguistic and educational standpoint, noting its advancements over previous rule-based systems. His study highlights ChatGPT's improved coherence and contextual relevance compared to earlier assistants like Siri and Alexa. However, Rižnar stresses the necessity for critical engagement with AI outputs, warning against uncritical acceptance of their authority in educational contexts.

Kai-Fu Lee and Chen Qiufan's *AI 2041* (2021) adopts a speculative yet analytical approach to envision AI's societal transformation by 2041. Through fictional narratives followed by contextual essays, they emphasize both the potential benefits and risks of AI, such as job displacement and privacy erosion. They recognize current AI's limitations but anticipate significant future advancements. Because it blends fiction with forecasts, *AI 2041* inherently involves a lot of speculation. Some scenarios may turn out to be unrealistic, too optimistic, or miss key developments that could change the landscape dramatically before 2041. Lee and Qiufan acknowledge risks but do not heavily critique the role of big tech companies in driving problematic AI deployments, which some critics see as a blind spot given Lee's background as a former executive at Microsoft, Apple, and Google China.

Jeff Hawkins' *A Thousand Brains* (2021) offers a critical theoretical shift by proposing the 'Thousand Brains Theory of Intelligence'. Hawkins critiques current AI systems like large language models for their lack of true understanding. He advocates for AI designs inspired by the brain's

decentralized model-building, emphasizing the importance of movement and spatial cognition for genuine intelligence.

Arvind Narayanan and Sayash Kapoor's *AI Snake Oil* (2024) provides a critical lens, warning that much of today's AI fails to deliver on its promises. They distinguish between generative and predictive AI, cautioning against trusting AI in high-stakes environments like healthcare or hiring without rigorous validation. The authors argue that many AI applications are overhyped and lack empirical validation, leading to potential harms when misapplied. While acknowledging the potential of some AI, such as ChatGPT, *AI Snake Oil* uncovers rampant misleading claims about the capabilities of AI and describes the serious harm AI is already causing in how it is being built, marketed, and used in areas such as education, medicine, hiring, banking, insurance, and criminal justice. The authors also discuss how academic research can contribute to AI hype, often due to flawed methodologies or overstated conclusions. They highlight the 'reproducibility crisis' in AI research, where many published results cannot be replicated or verified. This section emphasizes the need for more rigorous standards and transparency in AI research. Their call for widespread AI literacy reflects a broader concern about the misapplication of technological tools.

Meredith Broussard's works, *More Than a Glitch* (2023) and *Artificial Unintelligence* (2018), complement these critiques by illustrating how AI systems often perpetuate societal biases. Broussard challenges 'technochauvinism' – the blind faith in technological solutions – and advocates for algorithmic accountability, diverse development teams, and a balance between human judgment and computational power. Broussard calls for the implementation of algorithmic auditing and accountability measures. She underscores the importance of diverse teams in tech development and the need for transparency in algorithmic decision-making processes to mitigate embedded biases. In her earlier work (Broussard, 2018), she emphasizes that AI systems lack true understanding and are constrained by the quality and scope of their training data. She illustrates this with examples such as AI models failing to account for socioeconomic factors in educational assessments.

Similarly, Melanie Mitchell's *Artificial Intelligence: A Guide for Thinking Humans* (2019) emphasizes that today's leading AI systems, especially those based on deep learning, are very powerful

at narrow tasks – such as playing complex games (like Go), recognizing images, translating text, or even generating human-like writing. However, their “intelligence” is highly specialized and superficial. These systems do not truly *understand* the tasks they perform; they learn statistical patterns from massive datasets and apply them without genuine comprehension. For example, an image recognition system might correctly label pictures most of the time, but it can also be easily fooled by slight distortions that a human would have no problem ignoring. Mitchell points out that even the most advanced AI models lack common sense, a basic understanding of how the world works, which humans (even young children) naturally possess. In short, AI today can seem impressively intelligent, but only within very limited domains and without deeper reasoning abilities. Mitchell argues strongly that we must be very cautious about giving AI systems authority over important life decisions. One major problem she identifies is *bias*: AI models often inherit biases from the data they are trained on. For instance, if a hiring algorithm is trained on biased historical hiring data, it will replicate and even amplify those biases in its recommendations. Moreover, AI systems are often ‘black boxes’ – their decision-making processes are opaque even to their creators. This lack of transparency makes it difficult to predict or correct their errors. Mitchell also highlights that AI can be fragile: minor changes in input data can lead to wildly incorrect outputs. Therefore, while AI can be a helpful tool, it is risky to entrust it with critical decisions like medical diagnoses, criminal justice outcomes, or financial approvals without careful human oversight and accountability mechanisms. Mitchell is sceptical of the idea that AI will surpass human intelligence anytime soon. She critiques the ‘AI hype’ often found in media and some technology circles, arguing that creating truly general, human-like intelligence is vastly more difficult than many predictions suggest. The current successes of AI are mostly in *narrow AI*, not in building systems that can reason, plan, learn flexibly, and adapt to new situations like humans can. She notes that humans are equipped with a rich combination of intuition, bodily experience, emotional intelligence, and cultural understanding – dimensions that current AI lacks entirely. Moreover, many of the most basic features of human cognition, like understanding cause and effect or forming concepts based on few [Is the correct meaning ‘a few’?] ex-

amples, remain largely unsolved in AI research.

Gerd Gigerenzer’s *How to Stay Smart in a Smart World* (2022) underscores human superiority in navigating uncertainty. While algorithms excel in predictable environments, Gigerenzer argues that human heuristics outperform AI in complex, unpredictable real-world settings. Gigerenzer shows that humans use heuristics (mental shortcuts) effectively in messy, uncertain environments. Instead of overloading on data or trying to predict everything, people can make smart, adaptive decisions with limited information. He warns against overreliance on algorithms and stresses the importance of maintaining human oversight where adaptability and ethical judgment are required. In Gigerenzer’s (2022, 8) words: ‘Machines can beat us when the world is like a game of chess or Go, but they stumble when uncertainty rules. In such cases, human intuition and simple rules often work better than complex calculations.’

Bennett’s *A Brief History of Intelligence* (2023) traces the evolutionary development of human cognition to inform future AI. He identifies stages like steering, reinforcing, simulating, mentalizing, and speaking as critical to human intelligence. Bennett discusses how studying models like ChatGPT can provide insights into both the achievements and shortcomings of current AI systems, especially regarding genuine comprehension and adaptability.

Mark Coeckelbergh’s *AI Ethics* (2020) extends the discourse into the philosophical domain. He argues that AI ethics must move beyond technical checklists to address systemic societal structures, political power, and human–machine relationships. Coeckelbergh emphasizes participatory ethics, where affected communities have a voice in shaping AI development. He highlights the following ethical issues: (1) Bias and injustice – AI often amplifies existing societal biases unless critically examined and corrected at a systemic level; (2) Responsibility and accountability – distributed systems like AI challenge traditional notions of individual responsibility, which requires new models of collective responsibility; (3) Opacity and explainability – as AI decisions become harder to interpret, there is a moral duty to demand transparency and understandable explanations; (4) Automation and work – AI’s impact on employment raises ethical concerns about dignity, purpose, and the right to meaningful labour; and (5) Surveillance and control – data collection and monitoring by AI systems threaten privacy

and individual freedoms.

Roberto Pieraccini's *AI Assistants* (2021) offers a pragmatic look at virtual assistants like Siri and Alexa. While Pieraccini provides an accessible overview of their technological foundations and future prospects, there is a clear tendency in his writing toward optimism and limited engagement with ethical and societal concerns.

Collectively, this literature paints a nuanced picture of AI and AI assistants. While acknowledging remarkable technological progress, these authors consistently caution against technological determinism, emphasize the need for critical oversight, and call for a socially embedded approach to AI development.

The discourse surrounding artificial intelligence (AI), particularly large language models (LLMs) like ChatGPT, has intensified as these systems become increasingly integrated into various facets of society. While their capabilities in generating human-like text are impressive, a critical examination reveals significant limitations, especially concerning genuine understanding and reasoning. Noam Chomsky, along with Ian Roberts and Jeffrey Watumull, articulates a compelling critique of LLMs in their article, 'The False Promise of ChatGPT', published in *The New York Times*. They argue that models like ChatGPT operate without any comprehension of the content they produce. LLMs process and generate text based on patterns in data, devoid of any awareness or intent.

### Research Methodology and Research Limitations

In order to gain a better insight into the familiarity and use patterns of ChatGPT of students in a higher education institution in Slovenia (at the University of Primorska), the authors created a survey via the 1KA platform (<https://1ka.arnes.si/>). The survey comprised 15 questions, of which three were included to describe the demographics of our population sample, and 12 of which pertained to the activity patterns and knowledge of students and professors with ChatGPT.

The link to this questionnaire was then sent to several groups of students studying at the University of Primorska in Slovenia. The questionnaire was open from the 29th of January 2024 until the 27th of February 2024.

The authors relied on the goodwill of students to complete the survey, which restricted the number of respondents. Furthermore, since all of the authors were based in or around Koper,

Slovenia, it was decided to use students of the University of Primorska as a sample population. It is important to note that the results may not represent the students of other universities in Slovenia. The reliability of the data was also imperfect in the sense that it rests upon a subjective assessment made by the respondents in an uncontrolled setting.

Because the survey was created with little anticipation of the results received, a second questionnaire could be more relevant and a refined set of questions might prove to be more insightful.

### Results

The questionnaire was opened by 290 students of the University of Primorska, of whom only 90 students answered.

Gender distribution among the respondents revealed that 35.56% were male, and 64.44% were female.

Among all the faculties at the University of Primorska, the Faculty of Management had 48 students who responded to the questionnaire. The Faculty of Mathematics, Natural Sciences, and Information Technologies had the highest response rate, with 49 students participating. Conversely, the Faculty of Humanities had only three respondents. Regrettably, there was no participation from the Faculty of Education. We received only one response from the Faculty of

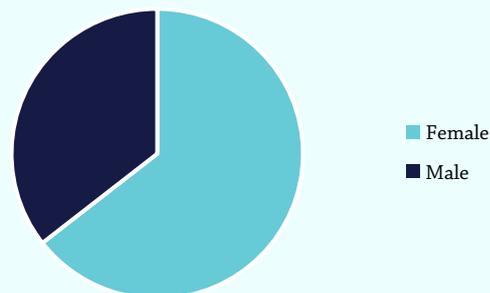


Figure 1 Gender distribution among the respondents

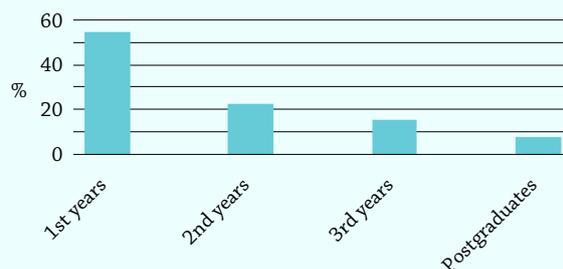


Figure 2 Population Demographics

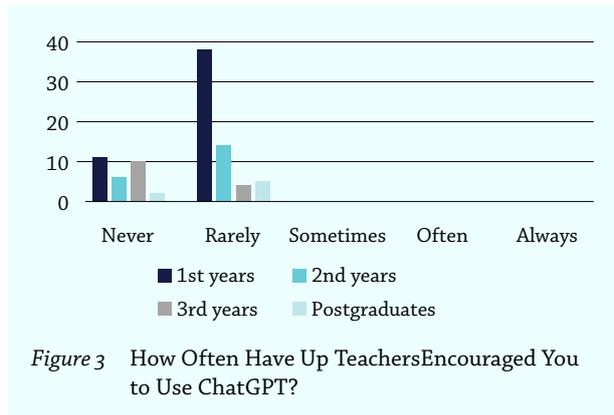


Figure 3 How Often Have Up Teachers Encouraged You to Use ChatGPT?

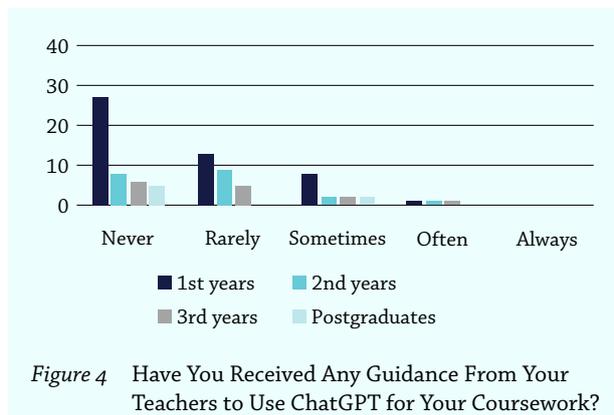


Figure 4 Have You Received Any Guidance From Your Teachers to Use ChatGPT for Your Coursework?

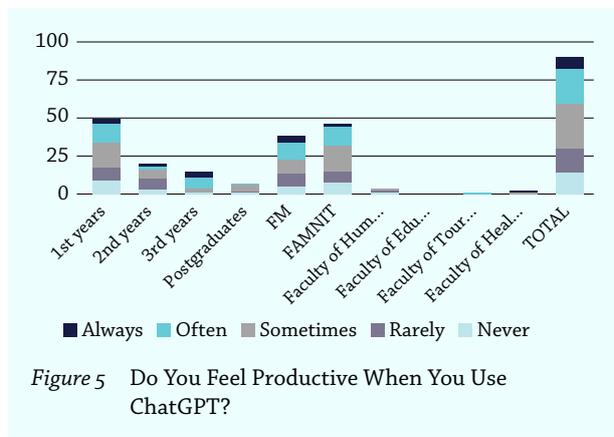


Figure 5 Do You Feel Productive When You Use ChatGPT?

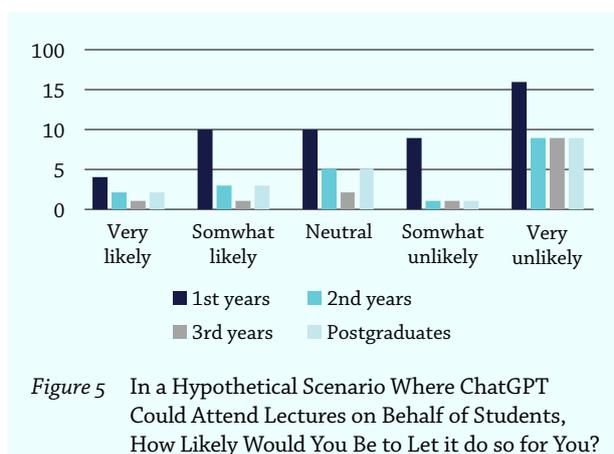


Figure 5 In a Hypothetical Scenario Where ChatGPT Could Attend Lectures on Behalf of Students, How Likely Would You Be to Let it do so for You?

Tourism, while the Faculty of Health Care provided two responses. Why students at these three faculties were not eager to participate is unclear.

Our survey on students’ utilization of ChatGPT for academic purposes revealed significant insights. The two main uses for ChatGPT were academic research (56.67%) and generating ideas (51.11%). Among the frequent use of ChatGPT, a considerable fraction admitted to never checking ChatGPT’s output (16.67%) or rarely checking (8.89%). A significant portion (32.22%) disclosed that professors never encourage the use of ChatGPT, while the majority (67.78%) receive rare endorsements, indicating a gap in institutional support.

Over half of the students (51.11%) never received guidance from their professors on its use, highlighting potential challenges in integrating ChatGPT effectively into educational activities.

The findings also indicate that 15.56% of respondents never feel productive when using ChatGPT. Lastly, 23.22% of respondents never or rarely find it useful to use ChatGPT, and 38.89% sometimes find it useful to use the bot.

In a hypothetical scenario, only a very small number of respondents would allow ChatGPT to attend lectures on their behalf.

### Future Research and Conclusion

Despite the above criticism, AI systems like ChatGPT have practical applications when used judiciously. They can assist in drafting content, summarizing information, and providing language support. However, it is crucial to recognize their limitations and ensure human oversight in their deployment. Hawkins’ work (2021, 113–117) challenges the success of current machine learning systems. While systems like large language models (ChatGPT-4, etc.) show impressive surface-level capabilities, they do not fundamentally *understand* in the way humans do. Hawkins argues (pp. 242–245) that real AI should be modelled after the architecture and operation of the human neocortex. By introducing the idea that multiple small, independent models working in parallel create robust intelligence, Hawkins suggests (pp. 124–128) a shift from monolithic AI architectures to more distributed ones. This is a call for ‘thousand brains’-style AI, where multiple semi-independent modules build their own models and then collaborate to create intelligent behaviour. Another important point is the emphasis on movement and sensory interaction. Unlike current disembodied AI systems that process static data,

Hawkins believes (pp. 145–170) true AI must move and interact with its environment to build accurate models – much like how a child learns.

The current landscape of AI, exemplified by systems like ChatGPT, presents a paradox. While these models demonstrate remarkable capabilities in language generation, they fundamentally lack understanding and reasoning. As Chomsky et al. (2023) state, the outputs of such systems are devoid of meaning, raising questions about their role and reliability in society; therefore, a critical and ethically informed approach is essential to navigate the promises and pitfalls of AI.

In Gigerenzer's (2022, 1–7) words, we need to be smart users of smart technologies – which means both knowing its limits and its powers; human intelligence thrives where unpredictability exists, and staying smart means understanding both what algorithms can do and what they cannot. In healthcare, for example, staying smart means using AI as a tool, not a master: supporting human decision-making without blindly outsourcing critical thinking.

The survey results underscore the pressing need for structured guidance and critical engagement with AI tools like ChatGPT within higher education. Although most students are aware of and utilize ChatGPT, their practices reveal concerning patterns of overreliance and insufficient critical verification of AI-generated content. The finding that 16.67% of respondents never check the output, and an additional 8.89% rarely do so, suggests a worrying gap in digital literacy and academic rigour. Furthermore, the minimal encouragement from professors – with over half of the students never having received any guidance on the use of ChatGPT – points to an institutional lag in adapting to emerging technologies. The fact that many students only sometimes or rarely find ChatGPT genuinely useful further highlights the mismatch between its potential and its actual educational value when used without appropriate framing or instruction. These insights collectively suggest that future efforts should not only aim to teach technical usage but also foster a culture of critical AI literacy – encouraging students to question, verify, and understand AI outputs rather than accept them passively. Higher education institutions should therefore prioritize integrating AI literacy into curricula, providing training for both students and faculty to ensure that AI tools are leveraged responsibly, ethically, and effectively to support – rather than undermine – learning and academic integrity.

## References

- Bennett, K. 2023. *A Brief History of Intelligence: Evolution, AI, and the Five Breakthroughs That Made Our Brains*. Harper.
- Broussard, M. 2018. *Artificial Unintelligence: How Computers Misunderstand the World*. MIT Press.
- Broussard, M. 2023. *More than a Glitch: Confronting Race, Gender, and Ability Bias in Tech*. MIT Press.
- Chomsky, N., I. Roberts, and J. Watumull. 2023. 'The False Promise of ChatGPT.' *New York Times*, 8. marec. <https://www.nytimes.com/2023/03/08/opinion/noam-chomsky-chatgpt-ai.html>.
- Coeckelbergh, M. 2020. *AI Ethics*. MIT Press.
- Farhat, F., S. Saquib Sohail, and D. Øivind Madsen. 'How Trustworthy Is ChatGPT? The Case of Bibliometric Analyses.' *Cogent Engineering* 10 (1): 2222988.
- Gigerenzer, G. 2022. *How to Stay Smart in a Smart World*. MIT Press.
- Hawkins, J. 2021. *A Thousand Brains: A New Theory of Intelligence*. Basic Books.
- Kelleher, J. D. 2019. *Deep Learning*. MIT Press.
- Lee, K.-F., and C. Qiufan. 2021. *AI 2041: Ten Visions for Our Future*. Currency.
- Mahuli, S. A., A. Rai, A. Vasant Mahuli, and A. Kumar. 2023. 'Application ChatGPT in Conducting Systematic Reviews and Meta-Analyses.' *British Dental Journal* 235 (2): 90–92.
- Mitchell, M. 2019. *Artificial Intelligence: A Guide for Thinking Humans*. Farrar, Straus and Giroux.
- Narayanan, A., and S. Kapoor. 2024. *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*. Princeton University Press.
- Pieraccini, R. 2021. *AI Assistants*. MIT Press.
- Rahman, M., H. J. Terano, N. Rahman, A. Salamzadeh, and S. Rahaman. 2023. 'ChatGPT and Academic Research: A Review and Recommendations Based on Practical Examples.' *Journal of Education, Management and Development Studies* 3 (1): 1–12.
- Rižnar, I. 2024. 'ChatGPT Usage: A Linguist's Perspective.' *Palimpsest* 9 (17): 191–200.