

LARGE LANGUAGE MODELS IN THE ENGINEERING WORKPLACE AND ASSOCIATED CURRICULAR IMPLICATIONS: AN EXPLORATORY STUDY

Jeroen Lievens

Faculty of Engineering Technology, University of Leuven (Belgium)

Abstract

The disruptive appearance of ChatGPT in November '22 has spurred considerable debate and research on student use of Large Language Models (LLMs) in higher education. Although divergent viewpoints exist, many institutions of higher education have gravitated towards a policy that can be succinctly characterized as “informed and responsible use”. A categorical ban on generative AI tools would fail to adequately prepare students for the future workplace where such tools are anticipated to be harnessed for their productivity enhancement, so the argument goes. Conversely, an all too uncritical embrace would undermine indispensable learning objectives of higher education curricula. This exploratory research paper seeks to identify the extent - and purposes - to which professional engineers are currently using LLMs and to examine associated curricular and pedagogical implications. In a student-driven action research project, engineering students of the second bachelor of the industrial engineering technology program at University of Leuven (Belgium) surveyed 249 engineers in October 2023. Results show that about half of respondents do not make use of LLMs for professional purposes and have no immediate intention of doing so, with some engineers referring to prohibitive corporate policies. About one-third of respondents are currently not employing LLMs, but they state the intention to explore their potential in the (near) future. The remaining respondents state that they already make use of LLMs for professional purposes, with over half using them for content generation. Other common modes of usage are instructing an LLM to revise a self-written text to optimize phrasing, spelling and grammar or to repurpose it for different audiences or media; to summarize texts; to write computer code; to explain technical concepts; to provide references or sources; and as a search engine. It is worth noting that several of these types of usage do not fit within the commonly accepted boundaries of “informed and responsible use”, underscoring the need for didactic interventions in higher education that raise student awareness of how LLMs actually function, what their inherent limitations are and which ethical concerns they entail. This paper describes how such interventions can be designed and integrated within an engineering program. Furthermore, it suggests ways in which higher education programs can monitor the fast-evolving landscape of AI workplace practices to ensure students are well-prepared to navigate the opportunities but also the challenges presented by LLMs.

Keywords: *Engineering education, large language models (LLMs), generative ai, ChatGPT.*

1. Introduction

Since the end of November 2022, teachers in higher education may have observed an unexpectedly high quality of academic writing in some student papers, alongside surprising substantive, factual, or logical errors. It was quickly understood that at least some students were happily outsourcing their tasks to ChatGPT, a generative AI (GenAI) application that is built on a pretrained large language model (LLM) (Cassidy, 2023). Such language models tend to be particularly proficient at writing, but being the “stochastic parrots” that they are, they are also prone to hallucinations and factual mistakes.

Responses from institutions for higher education were varied (Sullivan, 2023). Some universities argued for an outright ban of ChatGPT and other GenAI tools, and a return to invigilated examinations. The main reason cited was academic integrity: use of GenAI is understood to be paramount to plagiarism, as students are passing off work as their own while it is clearly not. Moreover, using GenAI, students are not meeting the learning objectives embedded in the task design; they are effectively bypassing the learning process and that bypass should be blocked off by clear regulations (Brody, 2023). Some universities, in contrast, were not so much focused on the potential threats of GenAI but rather on its

promise. Many commentators have pointed out that GenAI has the potential for revolutionizing the white-collar labor market much in the same way as the industrial revolution did for manual labor. If cognitive functions can be automatized, then the productivity of post-industrial economies is destined to accelerate, if not explode. Recent research has found that, indeed, the use of generative artificial intelligence substantially increases white-collar productivity while also compressing productivity distribution, as low-ability workers profit more from using GenAI tools (Noy & Zhang, 2023). From this viewpoint, it seems very imprudent - or even irresponsible - to educate students for their professional future while shielding them from a technology that is likely to transform the workplace in fundamental ways. Therefore, some universities have opted for a full embrace of this new technology rather than a unilateral ban. The university of Florida, for instance, has opted for an “AI Across the Curriculum” approach that integrates GenAI into all courses of the curriculum. The aim is to create an “AI-ready workforce” that is equipped to tackle the challenges of a rapidly changing workplace (Southworth et al. 2023).

In between these extremes, it appears that most universities are currently gravitating towards a more nuanced viewpoint that takes into account both the pitfalls and the promises of GenAI. Such universities argue for a responsible, informed and critical use of GenAI, allowing it within certain bounds. At University of Leuven (Belgium), for instance, students are allowed to use GenAI tools as a language assistant for checking or rewriting self-written text or as a search assistant to gather information, on the condition that the teacher did not provide instructions otherwise. Students are always expected to be transparent about their use of ChatGPT, and to acknowledge it clearly and reference it properly. Teachers, on the other hand, are expected to make explicit to students what the bounds are within which GenAI may or may not be used for certain tasks and also to make students aware of the limitations and pitfalls inherent in LLM technology.

This paper focuses on two research questions: a) which interventions can be integrated into course design to make students aware of the potential, but also of the problems linked to using LLMs? And b) to what extent, and how, do engineers in the workplace currently make use of Generative AI in their professional role? These questions will be explored in the context of a course in academic writing for second-year students in Engineering Technology at the University of Leuven, Belgium, at Campus Diepenbeek. Both questions share the same intention, which is to align the structure and shape of higher education onto the needs and realities of the workplace, both current and future.

2. Method

In the second-bachelor course Statistics+, we have been running an action research project with students and the workplace for more than ten years now. The main focus of the action research project is to investigate and monitor the perceived importance of several communication skills and tools for professional engineers. Students send out a standard questionnaire to engineers in their wider circle of acquaintance. Then, they process the quantitative data and supplement the resulting figures with qualitative data in the form of interviews with engineers. Finally, they write out their findings in an academic paper using a self-formulated research question. Statistics+ is a compound course, consisting out of sessions relating to statistics and quantitative data analysis as well as sessions dedicated to academic writing. Witnessing the sudden emergence of LLMs in the course of 2023, we decided to add a question to the standard survey focused on exploring the current use of GenAI tools by professionals. The 2023 survey was completed by 249 engineers in between September 29th and October 6th.

To answer the first research question, the teachers of the academic writing classes developed a range of small in-class interventions to make students aware of both the potential and the pitfalls of LLM-based GenAI tools. Some examples are listed here.

- Students were asked to write an argumentative paragraph prompted by the question: “Are communication skills important for engineers?”. Students were then asked to prompt this question to ChatGPT, and to compare their paragraph to the one outputted by ChatGPT. Students were then instructed to compare content and argumentation, the use of a topic thesis statement, the use of signal words, academic register and spelling and grammar correctness.
- Students were asked to prompt ChatGPT to add academic sources to the paragraph written by the tool to corroborate the proposed arguments. Students were then asked to verify the veracity of the proposed sources.
- Students were given an academic paper relevant to the topic. They were asked to write a synopsis of the paper, and then they were asked to prompt ChatGPT to write a synopsis as well. Then, students were asked to compare their own synopsis, the ChatGPT-generated synopsis and the abstract of the paper.

- Students were asked to research which guidelines are currently provided by academic journals regarding the use of ChatGPT.

A pre- and post-test regarding their use and their perception of GenAI tools was developed to analyze whether these interventions affected their understanding of and attitudes towards these tools. Since the post-test was not completed at the time of writing this paper, the findings of this pre- and post-test are not included in this paper.

To answer second research question, the data generated by the question that was added to the questionnaire was used. The question was phrased as follows: “Do you use GenAI tools (e.g., ChatGPT) for professional purposes?”, with the following options, of which the respondents could select more than one:

- No
- Not yet, but I intend to investigate their potential
- Yes, to generate content
- Yes, to optimize self-written text for structure and register
- Yes, to optimize self-written text for spelling and grammar
- Yes, in other ways: ...

To complement these quantitative data, qualitative data was gathered through semi-structured interviews with six engineers that had indicated in the questionnaire that they were already using GenAI tools for their jobs. Questions included, but were not limited to:

- Which tools do you use and to which purpose?
- Which advantages and/or pitfalls do you discern?
- How do you relate to ethical issues surrounding the use of GenAI tools (data privacy, ecological concerns, authorship issues...)?
- What is your company’s corporate policy regarding the use of GenAI tools?
- How do you expect the use of GenAI tools to evolve in your domain?
- Which recommendations do you have for educators and institutions teaching the engineers of tomorrow, with regard to GenAI tools?

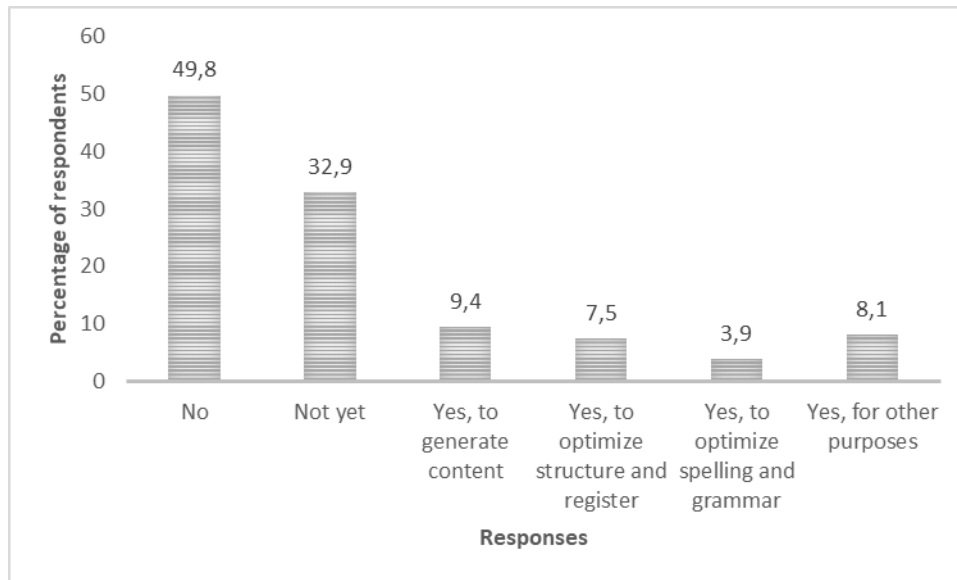
3. Results and discussion

From the in-class interventions, students learned that ChatGPT is proficient in matters of academic writing: the paragraphs produced by ChatGPT were led by a well-phrased topic thesis statement; the argumentation was solid and coherently structured using appropriate signal words; the register was academic; and linguistically, the paragraph was entirely correct. Comparing their own paragraphs to these machine-generated paragraphs, students could actually learn a lot, which highlights the didactic potential of LLMs in academic writing education. Students also learned that ChatGPT was quite able to generate a meaningful synopsis of an academic paper, which opens up an understanding of the ways in which such tools may be used to expediate the academic process.

However, when students verified the factuality of sources provided by ChatGPT, they were forced to recognize the disconnect between the LLM and the external, “real” world. ChatGPT obliged them with source references that closely simulated the structure and shape of actual references, but, which, upon closer inspection, turned out to be entirely fictitious. This finding occasioned insight into the operational design of such LLMs, which generate text on the basis of next-word statistical probability rather than real world meaning or correspondence. When investigating the guidelines issued by leading academic journals regarding the use of GenAI, students’ understanding of the capabilities as well as the limitations of LLM tools was further reinforced: Elsevier, for instance, allows the use of LLMs “to improve the readability and language of the research article, but not to replace key tasks that should be done by the authors, such as interpreting data or drawing scientific conclusions” (qtd. in Park, 2023, p. 105). Furthermore, the guidelines emphasize ethical use as they stipulate that authors must declare if and how they used an LLM-based tool, for instance in the Methods section of the paper.

For the second research question, quantitative data was obtained from the 249 engineers that completed the questionnaire in October 2023. Figure 1 shows the results.

Figure 1. Percentage of surveyed engineers making use of GenAI for professional purposes.



From these data, it appears that, in October 2023, close to 50% of the engineers responded that they did not make use of GenAI tools, with some referring to prohibitive company policies, while one-third of respondents indicated that they did not make use of GenAI yet, but that they intended to investigate the potential of such tools. Around 17% of responding engineers were already using GenAI applications in their jobs, with over half of them using the tool for content generation. Other common modes of usage are instructing an LLM to revise a self-written text to optimize phrasing, spelling and grammar or to repurpose it for different audiences or media; to summarize texts; to write computer code; to explain technical concepts; to provide references or sources; and as a search engine. Some of the above-listed functions are language-based, for which an LLM-based GenAI is indeed well-equipped, but other functions are clearly not. For instance, a LLM-based GenAI-tool is not suited for providing references or sources, as one of the in-class intervention made apparent. Also, an LLM-based GenAI is not a search engine; it is a conversational model trained on data but it does not have direct access to a search index. As such, it can be inaccurate and it is known to hallucinate. Using GenAI tools in professional contexts for purposes they are unfit for may lead to a host of unintended consequences, underscoring the need for proper training, whether in higher education or on the work floor.

The numbers produced by the small-scale questionnaire conducted by the second-bachelor students align quite well with findings published by Salesforce on the basis of a large-scale survey conducted in the same period (18-31 October, 2023) with over 14,000 employees across 14 countries. This survey revealed that 28% of employees were already using GenAI tools, a number that is slightly higher than our small-scale questionnaire produced. Quite similar to our findings, the Salesforce survey found that an additional 32% were expecting to use GenAI tools soon. The Salesforce survey also revealed that over half of the employees using GenAI tools were doing so without the formal approval of their employers. About 7 in 10 workers never received any training on how to use AI safely and ethically at work and a large majority of companies did not have clearly defined GenAI policies. These findings were echoed in the six semi-structured interviews conducted as part of our qualitative approach. Only two of the six interviewed engineers indicated that their employers had issued a GenAI policy stipulating use restrictions, while all of the interviewees expressed concerns relating to data privacy. Furthermore, it appeared that none of the interviewed engineers had been offered to engage in any training relating to GenAI. At the same time, all interviewees expected the significance of GenAI tools for job execution to increase in the future. In all, the interviews affirm the conclusion of the Salesforce survey that on the managerial level, efforts relating to GenAI training and guidelines are lagging behind a pioneering group of users.

4. Conclusion

Generative Artificial Intelligence defines an emerging field of workplace practices and use cases that is evolving at a rapid pace, and which needs to be monitored closely. Workforce adoption is picking up speed, while managerial concern seems to be trailing behind. This discrepancy underscores the importance of a comprehensive and well-designed approach to addressing the potential as well as the limitations of GenAI. Interventions such as described in the paper could be usefully integrated across the curriculum, highlighting the productivity increase that can be gained from these tools while also drawing attention to concerns relating to ethics, privacy, safety and ecology. The alignment of educational practices with the dynamic landscape of GenAI use in the workplace necessitates a) a continuous awareness of the changing realities of the workplaces, which can be extracted from a variety of sources (alumni meetings and other platforms where higher education meets workplace representatives, ongoing student action research projects such as described here etc.), and b) an agile approach to curriculum development, one that is responsive to the rapid technological changes and the evolving needs of the engineering profession. This approach can ensure that future engineers are prepared to meet the challenges and seize the opportunities of a rapidly transforming professional landscape.

References

- Brody, E. (2023, February 8). AI Writing tools garner concern about academic integrity. *Vermont Cynic*. Retrieved March 30, 2023, from <https://vtcynic.com/news/ai-writing-tools-garner-concern-about-academic-integrity-education-from-faculty/>
- Cassidy, C. (2023, January 17). Lecturers detects bot-use in one fifth of assessments as concerns mount over AI in exams. *The Guardian*. Retrieved March 30, 2023, from <https://www.theguardian.com/australia-news/2023/jan/17/lecturer-detects-bot-use-in-one-fifth-of-assessments-as-concerns-mount-over-ai-in-exams>
- Noy, S., & Zhang, W. (2023). Experimental evidence of the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187-192.
- Park, J. (2023). Could ChatGPT help you write your next scientific paper? Concerns on research ethics related to usage of artificial intelligence tools. *Journal of the Korean Association of Oral and Maxillfacial Surgeons*, 49(3), 105-106.
- SalesForce (2023, November 15). *More than half of generative AI adopters use unapproved tools at work*. Retrieved February 6, 2024, from <https://www.salesforce.com/news/stories/ai-at-work-research/>
- Southworth, J., Migliaccio, K., Glover, J., Glover, J.'N., Reed, D., McCarty, C., Brendemuhl, J., & Thomas, A. (2023). Developing a model for *AI Across the curriculum*: Transforming the higher education landscape via innovation in AI literacy. *Computers and Education: Artificial Intelligence*, 4, 100127.
- Sullivan, M, Kelly, A., McLaughlan, P. (2023). ChatGPT in higher education: Considerations for academic integrity and student learning. *Journal of Applied Learning & Teaching*. 6(1), 1-10.