ŌTAUTAHI CHRISTCHURCH AOTEAROA NEW ZEALAND
35TH AUSTRALASIAN ASSOCIATION FOR ENGINEERING EDUCATION ANNUAL CONFERENCE
8-11 DEC 2024
THE ENGINEER AND THE WORLD

Case study on generative AI Instruction in Civil Engineering Education at the University of Canterbury

Dempsey, David[a†]; van der Walt, Daniel[a†]; Nesbit, Trevor[b]; and Charters, Frances[a].
Civil & Natural Resource Engineering, University of Canterbury[a], Computer Science & Software Engineering, University of Canterbury[b]. These authors contributed equally[†]
Corresponding Author Email: david.dempsey@cantebury.ac.nz,

ABSTRACT

CONTEXT

The rapid adoption of generative AI (Gen AI) poses challenges for civil engineering education and its responsibility to teach foundational scientific knowledge, relevant technical skills, and responsible attitudes. Recent work has largely focused on challenges and opportunities for pedagogy, e.g., identifying and mitigating academic misconduct, and using Gen AI to augment instruction. Comparatively less attention has been paid to what knowledge, skills and attitudes graduate engineers require to be ready for workplaces with Gen AI.

GOAL

The goal of this study was to explore what effective instruction in Gen AI might look like for civil engineering students. More specifically, we wanted to know how students experienced a piloted 1-week instructional model, and what feedback they had for improvement. This study was undertaken so that we could improve future instruction on the topic in the department and faculty.

METHODOLOGY

A cohort of ten undergraduate civil engineering students were provided a 1-week instructional module on Gen AI. This was followed by a semi-structured focus group interview with five students that solicited their impressions and attitudes towards Gen AI, and their feedback on the instruction. Interview transcripts were thematically analysed with reflections presented.

ACTUAL OUTCOMES

Student feedback revealed an appreciation of the limitations of Gen AI, particularly as it relates to civil engineering practice. They also demonstrated some understanding of latent biases in language models. There was a discussion of the future of work, both negative (potential displacement of jobs) and constructive (identifying Gen AI use cases).

RECOMMENDATIONS

Based on the student feedback, Gen AI instruction can be strengthened. First, work-readiness can be improved by instructing on likely use-cases, including document query, Gen AI assisted coding, advanced prompt engineering. Second, increasing the length of instruction to strengthen particular skills, e.g., identification of misinformation, prompt engineering, and complex problem solving. Finally, student feedback should continue to be sought as instruction on this topic remains at an early stage, while Gen AI and its affordances continue to evolve rapidly.

KEYWORDS

Generative AI, AI Literacy, Gen AI Instruction, Student Feedback

Introduction

Undergraduate-level civil engineering education includes instruction across a range of competencies, including technical knowledge and skills, principles of engineering design, and the broader elements of engineering practice (Ning, 1996). For example, a graduate civil engineer might need to calculate the structural loads within a building, rank concept-level designs for wastewater treatment, or assess the social, sustainability, and financial characteristics of a new road network (International Engineering Alliance, 2021). The challenge for civil engineering educators is to provide effective instruction, exemplify the elements of good practice, and verify the attainment of learning by graduates.

The rapid emergence of generative AI (Gen AI) poses certain challenges and opportunities for civil engineering education. New methods for verifying learning outcomes will be needed given convincingly realistic generative text (Scarfe et al., 2024; Adilov et al., 2024; Nikolic et al., 2023). At the same time, educators have identified opportunities for Gen AI to improve learning experiences (Labadze et al., 2023; Baidoo-Anu & Ansah, 2023; Sanchez-Ruiz et al., 2023). Finally, much remains unknown about long-term pedagogical harms of shallow learning, learner dependency on AI (Bastani et al., 2024) and misinformation (Qadir, 2023).

Given the scale and pace of these challenges, it can be easy to overlook our key educational obligation: to provide good instruction on Gen AI, thus creating graduates that can use AI effectively while engaging in critical-thinking and problem solving. Given its strengths in text generation, modification, summarisation, and emergent competencies in technical calculation and few-shot prompting (demonstration learning; Liu et al., 2023), Gen AI is likely to find a range of use-cases in civil engineering practice, e.g., researching new topics, analysing datasets, ideation and evaluation during concept design, production of written reports, and assistance with peer review.

At the same time, there are concerns that an engineer's technical competency may be substituted by AI output of mixed reliability, as well as more widely discussed issues of data privacy and AI transparency. Thus, civil engineering educators should consider what Gen AI knowledge, skills and attitudes are important for their students. Defining this curriculum, developing effective learning artefacts, and then verifying learning outcomes are key tasks for building Gen AI literacy in civil engineering students.

Broad Aim and Goal

In this contribution, we present a case study and critical appraisal of a Gen AI teaching module in undergraduate-level civil engineering education. The goal of that module was to build Gen AI literacy specifically tailored to a civil engineering context. The case study describes the curriculum scope, learning activities, and assessments. Finally, we use focus group interviews to understand student experiences on the course and consider options to improve future offerings of AI instruction. For this study, we have adopted the Long and Magerko (2020) definition of AI literacy, as "*...a set of competencies that enables individuals to critically evaluate AI technologies; communicate and collaborate effectively with AI; and use AI as a tool online, at home, and in the workplace."* This definition is useful for engineers as it speaks to several graduate attributes of the internationally-agreed Washington Accord (International Engineering Alliance, 2021), notably Modern Tool Usage, Ethics and Communication. Finally, although Long and Magerko (2020) refer to AI technologies broadly, here we use the definition (without modification) to refer to Gen AI specifically.

Literature Review

Increasingly, Artificial Intelligence has been used as an umbrella term to refer widely to topics in natural language processing, machine learning, deep learning and intelligent agents. Although our focus here is on the comparatively narrow subfield of Gen AI, there is value in reviewing pedagogical literature drawn from the wider AI field for relevant attributes and practices.

 Copyright © David Dempsey, Daniel van der Walt, Trevor Nesbit, & Frances Charters, 2024

In their review of prior studies, Ng et al. (2021) explore how AI literacy (broadly defined) has been scoped by previous researchers, and catalog the strategies educators have taken to develop and evaluate AI literacy. The surveyed approaches to AI literacy were framed according to a 3-level Bloom's taxonomy: (i) know & understand; (ii) use & apply; and (iii) evaluate & create. Concurrently, the authors identified and categorised learning artefacts into development-related (hardware/software-focused), use-related (applying AI systems without having to build them), and unplugged (e.g., case studies, role-play). This framing aligns at a high-level with the components of computational thinking as outlined by Brennan & Resnick (2012) in discussing distinctive competencies of computer programming: (i) understanding computational concepts, (ii) applying computational practices, and (iii) appreciating computational perspectives. Ng et al. (2021) suggest a similar framework could be adapted to AI literacy broadly. We expect the framework may apply with few modifications to the subfield of Gen AI, presented in Table 1.

Table 1: Pedagogical frameworks in the context of Gen AI instruction. Examples are given in italics.

	Level 1	Level 2	Level 3
Bloom's Taxonomy (Bloom, 1956)	Know & understand. *Define Gen AI terms.*	Use & apply- *Write an effective prompt.*	Evaluate & create. *Verify Gen AI output for misinformation.*
Learning Artefacts (Ng et al., 2021)	Tool development- *Deploying and tuning an LLM.*	Tool use. *Use-ChatGPT for research and writing.*	Unplugged. *Undertake an organizational risk analysis for Gen AI.*
Gen AI thinking (Brennan & Resnick, 2012; Ng et al., 2019)	Understanding Gen AI concepts. *Language models,Misinformation*	Applying Gen AI practices. *Prompting and Use Cases.*	Critical assessment of Gen AI perspectives. *AI Alignment.*

The categories of Gen AI thinking map reasonably well to broader categories of engineering curricula that emphasise knowledge (e.g., mathematics science, economics), skills (e.g., problem solving, communication, research) and attitudes (e.g., society, environment) (Ning, 1995).

Research Gap

While significant strides have been made in the pedagogy of artificial intelligence a clear research gap remains in understanding the most effective ways to teach Gen AI to students in engineering education. Current pedagogical approaches often focus on technical proficiency and algorithmic understanding, however, there has been little discussion on what the learning outcomes should be for Engineering students. Additionally, it is unclear whether students require formal training as opposed to learning through passive use of technology.

Should the focus be on a student's ability to effectively generate novel content or rather on understanding the underlying principles that guide the AI's generative capabilities? What expectations around good practice and ethical behaviour do we set for Gen AI specifically, beyond the general expectations of engineering practice? Furthermore, how can these learning outcomes be effectively measured and assessed? These are questions that largely unanswered in the current body of research, indicating a gap to be addressed through further study.

This lack of knowledge hinders the development of intentional teaching strategies. It also limits our ability to provision graduate engineers with the skills and confidence to leverage current and future capabilities of Gen AI. Therefore, future research should aim to explore these areas, developing clear, evidence-based strategies for teaching Gen AI and establishing well-defined, meaningful learning outcomes.

Research Questions and Objectives

To address the research gap, this study investigates the following questions. What are the effective ways to teach Gen AI – specifically, what should the learning outcomes be for a Gen AI module? What are the students' perceptions of Gen AI instruction?

 Copyright © David Dempsey, Daniel van der Walt, Trevor Nesbit, & Frances Charters, 2024

It follows that the objectives of this study are to: 1) Present feedback from learners who have received Gen AI instruction, focusing on the effectiveness of the instruction and their perception of the important learning outcomes, and 2) Consolidate lessons learned from student feedback to inform future instruction and delivery of Gen AI.

Addressing these objectives will allow us to improve future curricula before providing it to a wider student cohort. It will also provide an evidence basis for further discussion around what the learning objectives should be when it comes to teaching Gen AI in an engineering context.

Methodology

ENCN404 is an elective course available to students studying the Bachelor of Engineering (Honours) in the Civil Engineering Department. The course was introduced as a special topic in 2024 to evaluate the long-term feasibility of instruction in modern modelling practices, including physics-based modelling, machine learning, and Gen AI.

Within the 12-week course curriculum, one week is dedicated to instruction on Gen AI. Three Learning Outcomes (LOs) were defined for this component of the course: 1) Students can define and understand components of Gen AI workflows, 2) Students can intentionally and effectively use Gen AI tools, and 3) Students can critically evaluate the risks of using Gen AI in professional engineering situations. These cover the three major components of AI-thinking (Table 1; Ng et al. 2019), while aligning with the traditional knowledge/skills/attitudes framing of engineering curricula (Ning, 1995).

Summative assessment of the LOs was via invigilated examination, which required students to demonstrate mastery at several levels of Bloom's Taxonomy. Formative assessments included a computer lab submission and a set of self-assessed comprehension questions provided at the end of lectures.

Course instruction was delivered via two 1-hour in-person lectures involving a mix of didactic instruction and student activities. Content was focused on standard Gen AI workflows, e.g., language models, training data, prompting, alignment, misinformation & verification, and ethics. Student activities during lectures were designed to illustrate major prompting competencies, e.g., providing specific details to the AI, iterating to improve prompts, and few-shot prompting. For these activities, students were directed to use the Bing Copilot on laptops they had brought to class. This tool was selected as students were able to obtain free access to GPT-4, classified as a frontier model at the time of instruction in early 2024.

Lecture concepts were reinforced during a 2-hour computer-based lab during which students used the same AI tool to complete a series of curated exercises at their own pace. These exercises were selected as either common Gen AI use cases, e.g., text editing, research, and writing, or because they illustrated some underlying property of the large language model, e.g., insensitivity to noise, encoded stereotypes. After each lab exercise, students were required to answer a set of reflection-style questions and record these in a document submitted at the end of the lab.

The final exercise in the lab required students to use the Gen AI for an engineering task: building a Risk Register and Safety Data Sheet for the hazardous chemical Ammonia. For the submission, they were required to reflect on which parts of the AI output would require technical verification (e.g., identified hazards and the completeness of the register) and which parts could be classified as an engineer's value judgement (e.g., assessment of hazard level, likelihood and consequences).

The final Exam of the course also included questions related to Gen AI, with a focus on definitions, understanding training data issues, prompt engineering and engineering ethics. Examples can be given on request. Course content was compiled from online sources, including open courses (Stanford University, CS324), discussion articles, and published literature.

 Copyright © David Dempsey, Daniel van der Walt, Trevor Nesbit, & Frances Charters, 2024

3.1 Study participants

The inaugural 2024 cohort of ENCN404 included 10 enrolled undergraduates and several PhD students who audited the course but did not participate in assessments. The interviewed students were undergraduates who completed all assessments. About three weeks after Gen AI instruction and before the final exam, undergraduates were invited to a focus group interview via the course management system, with an information sheet and consent form. This study received Human Ethics approval from the University of Canterbury Human Ethics Committee.

Semi-structured interview and analysis methodology

Students were introduced to a scenario involving the planning of a new subdivision and tasked with producing a report on the design of Three Waters, ensuring technical robustness, regulatory compliance, and client preferences. They discussed the possible use of Gen AI in this work. If not raised, the interviewer prompted discussions on verification and ethics. Researchers also asked students to reflect on how their views on Gen AI changed before and after the course and to provide recommendations for teaching this topic in engineering education. Basic thematic analysis was conducted on anonymised transcripts from a paid version of Otter.ai. The process began with familiarisation, where transcripts were read thoroughly. The AI-generated transcription was manually checked for accuracy. The transcript was then labelled by categorising parts based on content and meaning. Themes were developed to capture the essence of the data and reflect research objectives. Relationships, variations, and contradictions were explored, and findings were interpreted in the broader research context. Finally, the findings were presented clearly, using quotes to support key themes. This manual method did not use Gen AI but the manual work of researchers.

Results and discussion

The research findings are summarised by broad themes identified from the transcribed data, generally following the researchers' questions. Five students, representing ~20% of the class or 50% of enrolled students, participated in a focus group interview lasting just over an hour. Despite the small sample size, it includes a significant portion of the initial 10 enrolled students, who were typically higher-performing within the larger 190+ civil engineering cohort. Nevertheless, larger samples would be needed to confirm findings or establish statistically significant effects. Further, the self-selected nature of the group could create an inherent bias.

Engineering practice - The student's reflection on ways they would use AI in a project once in industry.

Students articulated awareness of the diversity of roles that Gen AI could have in supporting their work. However, a common theme was the importance of a user's engineering expertise to verify the AI outputs and the professional responsibility that remains with the engineer.

> *"Because we are engineers, and we have a lot of responsibility we will need to know how to use it responsibly".... "I will really refrain from taking content from ChatGPT . The prompt I would use here for my work, would be the style I am aiming for, it's a professional document please edit this text for coherency but do not change the meaning in any kind of way. "*

These comments recognize the dangers of uncritically copy-pasting Gen AI content in a professional setting. Instead, they suggest using Gen AI as a Copilot of sorts, which aligns with ENCN404 instruction on persona/role prompting (Schulhoff et al., 2024) and an emphasis that Gen AI act as an editor rather than as an original author.

> *"We are getting to the point now where we can upload massive, massive documents, another application that we can have now and moving forward, you can upload your entire literature, your internal documents your company has, obviously there are privacy concerns with that in itself, but you could upload your entire library, and say here are the standards that we have from the country and here is the companies information and what we care about - what are we missing, design something."*

 Copyright © David Dempsey, Daniel van der Walt, Trevor Nesbit, & Frances Charters, 2024

The student comment highlights current data management capabilities versus future applications. They describe use cases beyond report preparation, such as using AI for peer review or oversight ("what are we missing?") and direct substitution of engineering tasks ("design something"). These use cases were discussed during the course.

> *"You will have to rely on AI more and more to stay competitive in the market, it will be the professional's responsibility to take that output and make sure that it is in line with their expertise"*

This response also suggests the student is speculating not just on contemporary capabilities and uses, but also what the duties of a professional engineer might look like in coming years, assuming Gen AI becomes more widely adopted.

> *"I see yourselves falling into the same role as pilots do. An aeroplane can effectively fly itself, the pilot is there for responsibility and there to make the final call."*

Here, the student is anticipating the automation of certain engineering tasks through Gen AI. However, they are also connecting that delegation with an implied responsibility on the engineer to ensure quality. The autopilot analogy is an interesting one for the student to have used, and is not one that was used during the course. It is noteworthy that civil aviation authorities have at times expressed concern about pilot overreliance on automated systems and their detrimental impact on flying and aviation skills. A similar undermining of technical competency has been demonstrated for mathematics instruction when improperly augmented by Gen AI (Bastani et al., 2024).

Student reflection on ethical considerations with Gen AI

Use of Gen AI raises several ethical issues that need careful consideration by civil engineering graduates. Bias and fairness are concerns, as AI models can replicate biases present in their training data, potentially leading to unfair or discriminatory outcomes.

> *"There's always going to be a risk and a bias. From, you know, where you pull the information, and that's not just exclusive to ChatGPT, you know, you can ask a colleague or someone else as well for the input, and there'll be biases, obviously, [you need to] check."*

The student comment acknowledges bias, linking it to human sources rather than just training data. In engineering, Gen AI may advocate specific views, like design hold-points or hazard evaluations, reflecting professional opinion differences. The lack of transparency and accountability in AI decision-making challenges understanding and trust. Privacy concerns arise from using large datasets with personal information, and intellectual property issues emerge as AI-generated content may infringe on copyrights. The student comment above speculates in a rather free-wheeling way on these ethical considerations, highlighting the issue of whose ethics are prioritised within a model. Another student's comment reinforced this:

> *"There's a lot of talk in the industry between the different models. Open AI and Elon Musk has got one called Grok. And Elon Musk definitely has opinions on freedom of speech and all of that kind of stuff. And that's just within the Western world. There's massive discussions over how censored or open or free-thinking the AI should be. And then, you know, that's not even taking into account that a Chinese or Russian or Middle Eastern or an African AI models, they are going to have far different imposed rules, I guess."*

Student reflection on verification of Gen AI

A key pitfall of using Gen AI in engineering is over-reliance without sufficient human oversight. Gen AI might produce innovative designs or advise on complex processes but can introduce biases or errors compromising safety and functionality. Lack of transparency in AI decision-making prevents stakeholders from understanding risks and validating outcomes. Therefore, AI use must be balanced by rigorous human evaluation, relying on students' fundamental engineering skills and knowledge. The student below reflected some of these verification sentiments although not at a great level of depth. This is perhaps an area where deeper instruction is needed.

 Copyright © David Dempsey, Daniel van der Walt, Trevor Nesbit, & Frances Charters, 2024

"ChatGPT should only be an editing software, not an initial decision software, because it shouldn't be making that decision."

Student reflection on how their views have changed on their own use of Gen AI since receiving formal instruction in ENCN404.

Students were queried on what changes in their views, if any, arose from instruction on Gen AI. While this is a weak measure of the effectiveness of instruction, qualitative feedback is useful for reflecting on future course changes.

"Before [receiving instruction] I thought that I could rely on ChatGPT for more than my writing. But if I do become a professional engineer relying on ChatGPT can be very dangerous."

“More mindful of the limitations of each response [from Gen AI tools]."

"For engineering discussion, you may need to consider other aspects for example biculturalism that is not currently being addressed by AI."

“I Learned that ChatGPT can sound so convincing while being completely wrong.'"

"[Course instruction] showed us the limitations, the biases are important, it is very easy to not be aware of them, it does sound incredibly coherent and smart. When it does give you something that is false [or] biased it's convincing."

The comments above reflect that students have acquired some appreciation of the pitfalls of Gen AI and this has led to associated lowering of trust in the tools These comments reflect aspects of the instruction that emphasised the possibility of incomplete or subjective (acknowledging multiple viewpoints) outputs from Gen AI.

"...how to word your prompt to give the best output... depending on what you want... make it act like a tutor, you can make [it]... act like an editor, or make it like just proofread your work and not change the meaning.”

"Don't give things (from Instructor) that are easily 'ChatGPT'-able."

“So it might not actually just make us dumber, overall, but make us less proficient in some areas, but more proficient in others such so kinda just shifts our what we're good at."

One student highlighted the prompt engineering skills that they had acquired during the course. A student recognised the unresolved concerns around assessments and issues of academic misconduct, suggesting the difficulty for Gen AI to complete an assessment should be taken into consideration. Students recognize that there is a shift in skills and knowledge. Finally, this student's comment speculates on the potential for Gen AI to atrophy certain skills while developing new ones afforded by the technology. Similar arguments were made around the introduction of calculators and no doubt have been considered for other technological advances.

Student reflection on recommendations they would make to teach Gen AI in the context of Engineering education going forward.

This final question was intended to solicit the learner's perspective on Gen AI instruction, focusing on what they think they need in a future workplace.

"Suggestions from the university of what AI tool to adopt and where" and "where can Gen AI can connect to other Sofware like MathCAD."

These student comments discuss the affordances of different generative models, that is, what they do well and when they should be used. The student’s suggestion in this case is that instruction should be prioritized on the tools that will be most useful to them.

"It is the interaction that you have with your tutors and your lecturers. And there has, we've already seen a marked decline, because of technologies such as just recording everything, which is, you know, again, there are benefits to that. But it does mean that a subset of students will just check out and not interact with the lectures. There are things I get from you as well that I just don't get from your notes. One of the main risks is that students will [say] ‘you will never see me again, I will consult the AI and just watch the videos at home [and] submit the assignments online’... you lose your own learning."

 Copyright © David Dempsey, Daniel van der Walt, Trevor Nesbit, & Frances Charters, 2024

This student comment connects with other contemporary attendance issues on campus and their link to recorded instruction. In this case, they are expressing a concern Gen AI could replace human instructors entirely and this, ultimately, would be detrimental to their learning *“There are things I get from you as well that I just don't get from your notes.”*

> *"...you're going to need to teach students that [they] are replaceable if [they] rely on AI. So yes, you can go through your entire degree. And if you can get AI to do all your assignments for you, all the write-ups, give you all the ideas, do the coding for you everything, then congratulations, here's a degree, and you're replaced, you're completely replaceable by an AI... you need to use the AI to make you better. I think I don't think it's about being better than the AI, I think you need to use it as much to your ability to augment your abilities to be better, to be more efficient to be more productive, whatever. Because if you aren't...and you're getting AI to do all your studies and all your work, you're replaceable.”*

This comment reflects a broader concern about the future of work and employability of engineering graduates as AI workflows become increasingly adept at technical tasks (as measured by their ability to complete university-level assessments). The concern of job displacement through automation is not new, although arguably is new to more technically specialized disciplines like engineering. When students were asked specifically whether they thought Gen AI should be a skill acquired passively through ordinary instruction, or something acquired intentionally through formal training, one responded.

> *"I think students should receive formal training as soon as possible from the start and constantly refresh that material, you will have to keep up with AI. It will not surprise me if every university will have its own department focused on AI."*

Recommendations and Future Work

Several student comments reflected workplace issues or concerns, which makes sense for a cohort that is one semester away from finishing their studies and entering the workforce. In future, instruction will be revised to include tools and use cases relevant to industry. However, defining precise use cases is difficult given the newness of the technology and mixed uptake in professional practice. Training workshops with industry stakeholders have highlighted at least one priority, which is the use of Gen AI to write code (e.g., Python) for data analysis. Gen AI coding has perhaps been one of the earliest and most widely adopted use cases. For civil engineers, who might be more likely than other disciplines to doubt their coding abilities, access to Gen AI could lower barriers to entry. Thus, it is likely that we will need to consider instruction on Gen AI coding and tools (e.g. Copilot for Visual Studio Code) across the wider undergraduate program. A renewed focus on work-ready skills could also help to address the anxieties expressed by some around the future of work and the ease of finding employment.

A limitation of the Gen AI instruction in ENCN404 is that it only lasted one week. Although sufficient to introduce the breadth of the topic, we were not able to spend long periods of time with individual use cases like text editing, document querying and summary, or advanced prompting strategies. In the future, this might be addressed through hackathon-style assessments that challenge the students to address a particular engineering problem in considerable depth, e.g. peer review of a complex study, or analysis of a large dataset. Doing so would provide a more extended period of debugging the Gen AI to build skills, while also offering more complex and nuanced discussions of risk and reward.

Although student comments did identify different types of misinformation that could arise through Gen AI use, there was little discussion on verification approaches to mitigate these risks. This, too, perhaps reflects the relatively high-level, didactic discussion on this topic arising from the time-constrained instruction period. In the future, a longer form exercise that prompts identification of different misinformation types in a larger text may improve skill development and improved alertness to the risks.

 Copyright © David Dempsey, Daniel van der Walt, Trevor Nesbit, & Frances Charters, 2024

This study has been a useful first step, in highlighting student perspectives. Ongoing discussions at the department and university levels will determine the manner and level of Gen AI instruction. Further research at the department level should include pre and post-instruction interviews and an expanded sample size at various year levels. This would allow incremental development of instructional content based on student feedback.

References

Adilov, N., Cline, J. W., Hanke, H., Kauffman, K., Meneau, L., Resendez, E., Singh, S., Slaubaugh, M., & Suntornpithug, N. (2023). ChatGPT and the course vulnerability index. *Journal of Education for Business*, 1–8. https://doi.org/10.1080/08832323.2023.2260929

BaiDoo-Anu, D., & Owusu Ansah, L. (2023). Education in the Era of Generative Artificial Intelligence (AI): Understanding the Potential Benefits of ChatGPT in Promoting Teaching and Learning. *Journal of AI*, 7(1), 52–62. https://doi.org/10.61969/jai.1337500

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2024). *Generative AI Can Harm Learning*. https://doi.org/10.2139/ssrn.4895486

Brennan, K., & Resnick, M. (2012). New frameworks for studying and assessing the development of computational thinking. https://scratched.gse.harvard.edu/ct/files/AERA2012.pdf

Labadze, L., Grigolia, M., & Machaidze, L. (2023). Role of AI chatbots in education: Systematic literature review. *International Journal of Educational Technology in Higher Education*, 20(1), 56. https://doi.org/10.1186/s41239-023-00426-1

Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., & Neubig, G. (2023). Pre-train, Prompt, and Predict: A Systematic Survey of Prompting Methods in Natural Language Processing. *ACM Computing Surveys*, 55(9), 1–35. https://doi.org/10.1145/3560815

Long, D., & Magerko, B. (2020). What is AI Literacy? Competencies and Design Considerations. *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–16. https://doi.org/10.1145/3313831.3376727

Nikolic, S., Daniel, S., Haque, R., Belkina, M., Hassan, G. M., Grundy, S., Lyden, S., Neal, P., & Sandison, C. (2023). ChatGPT versus engineering education assessment: A multidisciplinary and multi-institutional benchmarking and analysis of this generative artificial intelligence tool to investigate assessment integrity. *European Journal of Engineering Education*, 48(4), 559–614. https://doi.org/10.1080/03043797.2023.2213169

Ning, C. C. (1995). Undergraduate Academic Programme: Planning, Development, Implantation and Evaluation. *International Journal of Engineering Education*, 11, 175-184.

Qadir, J. (2023). Engineering Education in the Era of ChatGPT: Promise and Pitfalls of Generative AI for Education. 2023 IEEE Global Engineering Education Conference (EDUCON), 1–9. https://doi.org/10.1109/EDUCON54358.2023.10125121

Sánchez-Ruiz, L. M., Moll-López, S., Nuñez-Pérez, A., Moraño-Fernández, J. A., & Vega-Fleitas, E. (2023). ChatGPT Challenges Blended Learning Methodologies in Engineering Education: A Case Study in Mathematics. *Applied Sciences*, 13(10), 6039. https://doi.org/10.3390/app13106039

Scarfe, P., Watcham, K., Clarke, A., & Roesch, E. (2024). A real-world test of artificial intelligence infiltration of a university examinations system: A “Turing Test” case study. PLOS ONE, 19(6), e0305354. https://doi.org/10.1371/journal.pone.0305354

Schulhoff, S., Ilie, M., Balepur, N., Kahadze, K., Liu, A., Si, C., Li, Y., Gupta, A., Han, H., Schulhoff, S., Dulepet, P. S., Vidyadhara, S., Ki, D., Agrawal, S., Pham, C., Kroiz, G., Li, F., Tao, H., Srivastava, A., … Resnik, P. (2024). The Prompt Report: A Systematic Survey of Prompting Techniques (Version 3). arXiv. https://doi.org/10.48550/ARXIV.2406.06608

Stanford University. (2022). CS324: Winter 2022 – Advances in Foundation Models. *Stanford University*. https://stanford-cs324.github.io/winter2022/

Copyright statement

Copyright © 2024 David Dempsey, Daniel van der Walt, Trevor Nesbit, & Frances Charters: The authors assign to the Australasian Association for Engineering Education (AAEE) and educational non-profit institutions a non-exclusive licence to use this document for personal use and in courses of instruction provided that the article is used in full and this copyright statement is reproduced. The authors also grant a non-exclusive licence to AAEE to publish this document in full on the World Wide Web (prime sites and mirrors), on Memory Sticks, and in printed form within the AAEE 2024 proceedings. Any other usage is prohibited without the express permission of the authors.