

Mapping out a research agenda for generative artificial intelligence in tertiary education

Jason M. Lodge

The University of Queensland

Kate Thompson

Queensland University of Technology

Linda Corrin

Deakin University

Generative artificial intelligence (AI) has taken the world by storm. In this editorial, we outline some of the key areas of tertiary education impacted by large language models and associated applications that will require re-thinking and research to address in the short to medium term. Given how rapidly generative AI developments are currently occurring, this editorial is speculative. Although there is a long history of research on AI in education, the current situation is both unprecedented and seemingly not something that the AI in education community fully predicted. We also outline the editorial position of AJET in regards to generative AI to assist authors using tools such as ChatGPT as any part of the research or writing process. This is a rapidly evolving space. We have attempted to provide some clarity in this editorial while acknowledging that we may need to revisit some or all of what we offer here in the weeks and months ahead.

Keywords: generative artificial intelligence, research, assessment

Introduction

The seemingly rapid emergence and explosive growth in the capacity and use of generative artificial intelligence (AI) has taken tertiary education by surprise. This perhaps should not have been the case given the long history of research in the use of AI in some subfields of education inquiry (see Chen et al., 2020). An active community of researchers has been exploring the implications and use of AI for decades (Chen et al., 2022). Topics of interest in this research have spanned assessment (e.g. Zawacki-Richter et al., 2019), learning and teaching through and with AI (e.g. Kuka et al., 2022), and the technical and ethical aspects of AI relevant to education (e.g. Nguyen et al., 2023). However, none of this previous work seemingly predicted that powerful tools such as ChatGPT would be released publicly and widely available over such a short timeframe. With this history in mind, we map out an agenda for research on these technologies in the educational technology community in the short to medium term and discuss the editorial position AJET is taking on the new tools, like ChatGPT. We do so being fully aware that developments are moving so rapidly that there have been calls for a pause on further development (Future of Life Institute, 2023).

The generative AI we discuss here mostly refers to applications built on large language models (LLMs) such as GPT-4, the model underpinning ChatGPT. While there are other kinds of generative AI based on images or video, for example, LLMs are at the core of the issues currently facing tertiary education. LLMs are machine learning algorithms built to work with enormous datasets to create predictions based on a stimulus of some kind. In the case of ChatGPT, the stimulus is in the form of a prompt from the user. At the time of writing, there is still much to be figured out in terms of the rigorous collection and dissemination of the findings of studies examining how these LLMs and the applications built on top of them are impacting tertiary education. The evolution of generative AI is currently a day-by-day prospect with hundreds of new applications emerging daily. The hype that surrounds these technologies lacks much evidence to support firm claims about the impact and utility of the tools for education practice or appropriate adaptation of policy. There are strongly worded and persuasively argued articles in the grey literature spanning a range of opinions from the idea that generative AI will have little impact through to

claims about it fundamentally changing humanity forever, suggesting we have reached a tipping point of sorts.

In writing this editorial, we reviewed the recent peer-reviewed research on AI in tertiary education several times (for overview, see special issue edited by Gasevic et al., 2023). It is evident that much of the scholarly discussion and debate from as recently as last year is now sorely outdated given the emergence of tools such as ChatGPT. As such, any agenda we can suggest at this stage can only be speculative. Inevitably, this editorial is perhaps one of the more subjective that we have written. However, we feel that this statement is warranted given the uniqueness of the situation in which tertiary institutions and our community find themselves. Unlike the emergence of other educational technology tools similarly surrounded by significant hype, such as during the peak of the Massive Open Online Course (MOOC) discussion in 2012, generative AI has quickly become a major issue for tertiary education policymakers, institutions and practitioners, with limited time for careful consideration.

Even in these early days of such advanced AI capabilities, there are some key areas we have identified with the potential to be seriously impacted by the emergence of these new technologies (see also, Gasevic et al., 2023). We may find that the next iterations of the many generative AI tools currently emerging will change the landscape rapidly, as has been the case so far in 2023. However, the following issues will continue to be areas of concern where researchers, practitioners, and policymakers will need to work together to design and conduct research to provide evidence-informed guidance about the impact of generative AI on learning, teaching, and leadership in tertiary education.

Critical research areas

The following section of the editorial includes core issues and questions that have emerged since the release of ChatGPT and subsequent generative AI-based tools. Their potential to change how we approach the design and facilitation of learning and teaching is great, and now is a salient opportunity for educators to reflect on the purposes and pedagogical structures that underlie tertiary education. These areas for future research are necessary to inform a balanced and evidence-informed discussion on how different stakeholders, from individuals to institutions, can harness the strengths and industry-aligned practices of generative AI use while maintaining the integrity and honesty of academic practices.

Sensemaking

Perhaps the most pressing task for many in tertiary education has been to make sense of what generative AI is, how the tools flooding the market are using this AI technology, how they have built up their “knowledge”, and the implications for stakeholders and practice in education. This will undoubtedly be an ongoing task as the technologies evolve. A useful introduction has been published by a team from UNESCO (Sabzalieva & Valentini, 2023) specifically focused on generative AI and ChatGPT for tertiary education. Despite this useful overview, much uncertainty remains about foundational topics such as what generative AI is and how it works. For example, the question of whether or not tools such as ChatGPT have emergent properties is the subject of ongoing research. Until we understand more about how generative AI works and the data on which it is based, the trustworthiness of the outputs will remain questionable, likely biased in multiple ways, and conceivably dangerous at times. In depth understanding of such fundamental issues is needed to be able to consider the real impact that generative AI will have on tertiary education.

While some of us in tertiary education have disciplinary backgrounds in computer science, cognitive sciences, and cognate areas that enable an understanding of developments in generative AI, many do not have this foundation. Research in areas such as explainable AI (e.g. Arrieta et al., 2020) applied to educators in tertiary education will be critical for informing people without a background in these disciplines. Better understanding of generative AI will remain a critical issue for educational technology researchers for the foreseeable future. Therefore, sensemaking will continue to be something that the community will need to engage in as these technologies continue to evolve.

Assessment integrity

The major focus for education policymakers upon the release of ChatGPT was concern over the likelihood of it being used by students to cheat. In many disciplines, assessment of learning is carried out by student production of an artefact (such as a laboratory report, essay, code, etc.) demonstrating their learning, aligned with particular criteria. A key issue with generative AI is that it allows students to produce artefacts by entering the assessment instructions into a tool such as ChatGPT, without going through the process of learning themselves. Generative AI such as ChatGPT does not work in the way that a calculator works. A calculator actually performs the calculations required in order to reach an answer; this is important. Generative AI does not perform calculations, it does not go through the learning, it does not engage in thinking. ChatGPT and similar tools make predictions; they guess.

It is important to recognise that the vast majority of students do not want to cheat and are invested in their learning, and that cheating and academic integrity are not new considerations in tertiary education (e.g. Dawson, 2021). However, the wide availability of generative AI has meant that the barriers to engaging in cheating behaviour (in terms of effort and risk) have been lowered significantly, and the ability to detect cheating has become significantly more difficult, if not impossible. It is too early to know for certain whether or not students will engage in some form of cheating using this technology and how widespread these practices are likely to be.

Generative AI tools are already being embedded within other applications such as word processing software or slideware. Tertiary educators and institutions will need to consider their assumptions about learning, and therefore cheating, within the broader consideration of our purpose in order to navigate the socio-technical ecology of higher education. At least part of this is to prepare students for a world in which AI is a core part of the tools used in most careers. The importance of the educational technology community of practitioners and researchers our community to engage in sensemaking in relation to generative AI is essential to considering these issues of assessment integrity and role of institutions in providing evidence of learning for students and future employers.

Assessment redesign

An important response to the availability of generative AI for students to use to complete assessment tasks is to redesign the assessment. In most disciplines, assessment designs include the production of artefacts such as laboratory reports, essays, or research reports. In many cases, these types of tasks are fairly easy to generate using AI tools, although the quality may require thorough review by the student if a high grade is sought. Indeed, some educators are encouraging students to use generative AI to help draft written assessment tasks, and then assessing how the students critique and improve the resulting outcome.

A broad set of options are emerging about how to redesign assessment in response to the availability of generative AI. Some options rely on either a return to assessment approaches such as in-person exams or presentations, or recommend a scaffolded series of tasks focusing on the demonstration of students' thinking processes relating to the development of artefacts. Others aim to exploit weaknesses in the current versions of the platforms relying on generative AI. However, these approaches ignore the fact that the weaknesses of the models underpinning ChatGPT and other tools are likely to improve in subsequent versions. ChatGPT can reproduce misconceptions widely held (such as 'neuromyths') and large language models of this kind have been (in)famously labelled 'stochastic parrots' because of this tendency (Bender et al., 2021). If, for example, an assessment task relies on this observation, there is no guarantee that students would not figure out prompts to get around the problem or that an update to the model would not stop it from reproducing that misconception. Assessment redesign will need to adjust to the ongoing improvement of these technologies, beyond exploiting the mistakes that ChatGPT produces.

There are also opportunities that generative AI tools provide that allow students to engage with additional assessment support and real-time formative feedback as part of their learning. For example, students

could use a tool such as ChatGPT to generate multiple sets of ‘practice questions’ to help in preparation for exams. As mentioned, students could enter their written work into ChatGPT and ask for feedback on aspects like the flow of an argument or grammatical structure. With any of these approaches it would need to be made clear to students that these models can provide responses that can be misleading, or simply wrong, so applying a critical lens to any outputs they receive is important. However, this process of checking the reliability of responses can also prompt learning and build vital skills for dealing with generative AI in other areas of life and future careers.

In all this discussion of possibilities and challenges for assessment, now is an important opportunity to consider the nature of learning and knowledge and how we prepare students for using the tools based on generative AI for work and life. The drive to redesign assessment in ways that consider alternative measures of learning creates many research opportunities for the educational community. While there may be uncertainty about how generative AI is trained and implemented, research into the outputs for different assessment scenarios can guide educators towards more reliable and effective assessment designs for the future. Research into how and why students use generative AI in the ways that they do for assessment, and not just applying an academic integrity lens, will also be vital going forward to help educators engage students by better understanding their motivations for learning.

Learning and teaching with AI

Through adaptive tutoring and in other ways, there is potential for the use of generative AI to personalise learning. Personalised and adaptive learning (Morze et al., 2021; Peng et al., 2019) have been proposed and used mostly in discipline areas such as engineering, languages, mathematics and sciences, and tend to focus on learning outcomes that are relatively straightforward to measure (Xie et al., 2019). The idea of personalisation of learning is contested, to say the least. It also remains to be seen how much of the personalisation of learning can be achieved through the use of generative AI. A series of questions that are of interest to us is the potential of generative AI for supporting students through co-regulation and socially-shared regulation of learning (see Järvelä et al., in press). As with other possibilities we describe here, we are only beginning to make sense of these possibilities and are far from robust practices for making co-regulation with AI effective for learning. There are, therefore, many questions emerging about how best to use AI to guide the processes of high quality learning.

The role that generative AI will play in learning tasks is also far from clear. It is not difficult to imagine that a sophisticated chatbot built on a large language model could be an effective adaptive tutor, of sorts. At a very basic level, a tool such as ChatGPT can take on an explanatory role for students to interrogate to better understand a concept. Due to the adaptive nature of the tool, this approach to using technology for learning is different to the use of search engines, where learners are given options and critical thinking skills are applied to synthesise information from a variety of sources (Knight & Littleton, 2015). When learners interact with ChatGPT, the information is already synthesised, without attribution to sources (unless specifically requested). Research about learning with and from tools built on generative AI will be of great importance in how they are included in formal learning environments, and how to still ensure learners develop analytic and critical thinking skills.

There has been much discussion and debate about the possibilities for using generative AI to assist with reducing workloads for teachers. There is some early evidence suggesting that ChatGPT can provide effective feedback (Dai et al., 2023). However, the disciplinary domain in which this early study was carried out is data science and, similar to research about adaptive and personalised learning in higher education, there are questions about how procedural (as opposed to conceptual) the task for which the feedback was provided is. Feedback has already been provided in procedural domains for decades, such as learning with flight simulators. How the example of feedback provision by generative AI in this study translates across disciplines and domains remains to be seen.

Generative AI has potential to contribute to educators’ practice in lesson planning and the generation of learning objects and activities. The tools to do this are very new at the point of writing this editorial, and, from our own exploration as teachers, it is apparent that lesson plans produced by ChatGPT need to be

reviewed by an experienced teacher to work effectively. The overall work of educational and instructional design seems to require human input. The generic lesson plans produced still require knowledge of the context of learning taking place. Given the emerging power of both the LLMs and the applications being built on top of these models, we expect there to be rapid development in this space. If these are successful, they will build on the fundamental notion that high quality learning is relational. Understanding how these technologies can be used to assist with teaching in tertiary education will, therefore, require careful consideration.

Ethics and AI

While we will not attempt a broad survey of the many emerging ethical and moral issues that have arisen as generative AI evolves, we do recognise here that there are some serious implications in this regard. Despite the amount of supervised and reinforcement learning that has gone into training these LLMs, they still exhibit a range of harmful biases. The notion of stochastic parrots we described earlier is also evident here. The internet is awash with misinformation, falsehoods and the worst kind of generalisations and stereotypes. Consequently, when a large dataset is generated with a scraping of the internet as a key component, biases, factually incorrect information, and prejudice of various kinds are going to be evident in any output generated from these data.

Beyond the numerous questions about what kind of data has been loaded into the models on which the tools are built, the privacy implications of the tools and the ethics related to using them in any way in tertiary education are also of concern. Nguyen and colleagues (2023) have provided some foundational principles for the use of AI in education. These principles allude to the reality here that every opportunity for using LLMs and associated applications also raises ethical issues. We expect that there will also continue to be a need for ongoing investigation of this wide range of ethical issues for the foreseeable future.

AJET's position on the use of generative AI in publications

While the rapid evolution of generative AI tools will provide a multitude of opportunities for research into their impact and future, there are important considerations that need to be made in respect to how researchers engage with such tools as part of the research and publication processes. In line with other journals internationally, it is our view that generative AI tools, such as ChatGPT, should not be treated as an author or cited as an agent responsible for intellectual property. Ultimately, ideas that appear in the journal need to be attributed to an agent, and we take that to be a human. This is an important point. In the unlikely event that something appears in the journal that is in some way unethical, immoral or illegal, as editors, we need to have clarity about who is responsible. While there are and will continue to be debates about the possible sentience of large language models, there is not enough evidence in our view that generative AI is a responsible agent.

As such, The AJET editorial team acknowledge the increase in availability and use of generative AI in academic research and set out the following rules in relation to how authors and editorial members should work with generative AI:

- Generative AI cannot be listed as an author on AJET publications. An author must be able to agree to and be accountable for the aspects of their authorship under the CRediT taxonomy, which an AI cannot do.
- Authors need to acknowledge the contribution made by generative AI tools to any aspects of the research published. In the acknowledgement section the authors should outline the specific tasks AI was used to complete, including (but not limited to) research design, data analyses, data visualisation, text creation/editing, etc.

- AJET Reviewers do not have permission to use generative AI to complete any reviews of AJET articles. Sharing articles under review with third party AI providers for this purpose may contravene authors' intellectual property rights to their work.

Setting a research agenda to help prepare students for an AI world

The rise of generative AI is poised to rapidly transform many industries, and tertiary education has a significant role to play in preparing students for a world where it exists and is impacting almost all aspects of society. The challenges and opportunities related to preparing students for this new world are vast, and it is crucial that educators, institutions and policymakers take a proactive approach to ensure that students are well-prepared for the future. We will be best placed to do this if we are collecting evidence to inform decision making. Which approaches should we be adopting? Is the structure and function of tertiary education sufficiently preparing students for this world? The role of educational researchers will be crucial to informing policy and practice over the coming years.

One of the key challenges of preparing students for a world where generative AI exists is ensuring that they have the necessary skills and knowledge to work alongside and with these technologies effectively. However, as we return to our first area of interest of sensemaking, at this point, we have not identified what this knowledge and skills are. Students must be prepared to navigate the ethical and social implications of generative AI, such as the impact on privacy and the potential for biased decision-making. It is apparent that students will need to develop skills necessary to work alongside these technologies to be competitive in the job market, and to learn with AI, adapting to the rapid turnover in new tools and ways of working. Continuing to monitor the alignment between tertiary education and the needs of employers will be vital for ensuring that these skills are what are needed.

A research agenda that focuses on the capabilities and applications of generative AI is needed to provide a solid evidence base for the appropriate and effective use of such technologies in tertiary education settings. From research that can inform learning and assessment design, to a greater understanding of the skills necessary for students to learn in order to engage with AI in the workplace, the rapid pace of change is driving our need to expand our knowledge quickly in order to respond to the constantly evolving situation. As always, AJET welcomes good quality work in this area for publication to contribute to the wider conversation and help the tertiary education community respond to, challenge, and evolve along with the technology.

Author contributions

Jason Lodge: Conceptualisation, Writing - original draft, Writing - review and editing; **Kate Thompson:** Writing - review and editing; **Linda Corrin:** Writing – review and editing.

Acknowledgements

ChatGPT was used to generate ideas for the writing of this editorial.

For our first issue of the year, we would like to formally thank the AJET Associate Editors and Copyeditors, without whose contribution AJET would not exist. In particular, we would like to thank Kayleen Wood for her many years as a copyeditor at AJET. Her dedication to detail and ensuring that the research that is published in AJET is of the highest quality has made a significant contribution to this publication. We wish her all the best in her new position and look forward to working with her as part of the ASCILITE community in the future.

We would also like to thank the AJET Management Committee member and the ASCILITE Executive who continue to provide critical support to the journal.

Lastly, we would like to acknowledge the generous work of AJET reviewers who serve as the foundation for a high-quality journal such as this.

Thank you all.

References

- Arrieta, A. B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., & Herrera, F. (2020). Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58, 82–115. <https://doi.org/10.1016/j.inffus.2019.12.012>
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency* (pp. 610-623).
- Chen, L., Chen, P., & Lin, Z. (2020). Artificial intelligence in education: A review. *IEEE Access*, 8, 75264–75278. <https://doi.org/10.1109/ACCESS.2020.2988510>
- Chen, X., Zou, D., Xie, H., Cheng, G., & Liu, C. (2022). Two decades of artificial intelligence in education. *Educational Technology & Society*, 25(1), 28-47.
- Dai, W., Lin, J., Jin, F., Li, T., Tsai, Y., Gasevic, D., & Chen, G. (2023). *Can large language models provide feedback to students? A case study on ChatGPT*. EdArVix Preprint. <https://doi.org/10.35542/osf.io/hcgzi>
- Dawson, P. (2021). *Defending assessment security in a digital world: Preventing e-cheating and supporting academic integrity in higher education*. Routledge.
- Future of Life Institute (2023). *Pause Giant AI Experiments: An Open Letter*. <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>
- Knight, S. J. G., & Littleton, K. (2015). Thinking, Interthinking, and Technological Tools. In: Wegerif, Rupert; Li, Li and Kaufman, James C. eds. *The Routledge International Handbook of Research on Teaching Thinking*. Routledge International Handbooks of Education. Abingdon, New York: Routledge, pp. 467–478.
- Kuka, L., Hörmann, C., Sabitzer, B. (2022). Teaching and Learning with AI in Higher Education: A Scoping Review. In: Auer, M.E., Pester, A., May, D. (eds) *Learning with Technologies and Technologies in Learning. Lecture Notes in Networks and Systems*, vol 456. Springer. https://doi.org/10.1007/978-3-031-04286-7_26
- Gasevic, D., Siemens, G., & Sadiq, S. (2023). Empowering learners for the age of artificial intelligence. *Computers and Education: Artificial Intelligence*, 100130. <https://doi.org/10.1016/j.caeai.2023.100130>
- Järvelä, S., Nguyen, A., & Hadwin, A. (in press). Human and artificial intelligence collaboration for socially shared regulation in learning. *British Journal of Educational Technology*.
- Morze, N., Varchenko-Trotsenko, L., Terletska, T., & Smyrnova-Trybulska, E. (2021). Implementation of adaptive learning at higher education institutions by means of Moodle LMS. *Journal of Physics: Conference Series*, 1840(1), 012062. <https://doi.org/10.1088/1742-6596/1840/1/012062>
- Nguyen, A., Ngo, H.N., Hong, Y., Dang, B., & Nguyen, B.T. (2023). Ethical principles for artificial intelligence in education. *Education and Information Technologies*, 28, 4221 - 4241 . <https://doi.org/10.1007/s10639-022-11316-w>
- Peng, H., Ma, S., & Spector, J. M. (2019). Personalized adaptive learning: An emerging pedagogical approach enabled by a smart learning environment. *Smart Learning Environments*, 6(1), 9. <https://doi.org/10.1186/s40561-019-0089-y>
- Sabzalieva, E. & Valentini, A. (2023). *ChatGPT and Artificial Intelligence in Higher Education: Quick Start Guide*. UNESCO.
- Xie, H., Chu, H.-C., Hwang, G.-J., & Wang, C.-C. (2019). Trends and development in technology-enhanced adaptive/personalized learning: A systematic review of journal publications from 2007 to 2017. *Computers & Education*, 140, 103599. <https://doi.org/10.1016/j.compedu.2019.103599>

Zawacki-Richter, O., Marín, V. I., Bond, M., & Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education—where are the educators?. *International Journal of Educational Technology in Higher Education*, 16 (39), <https://doi.org/10.1186/s41239-019-0171-0>.

Corresponding author: Jason M. Lodge, jason.lodge@uq.edu.au

Copyright: Articles published in the *Australasian Journal of Educational Technology* (AJET) are available under Creative Commons Attribution Non-Commercial No Derivatives Licence (CC BY-NC-ND 4.0). Authors retain copyright in their work and grant AJET right of first publication under CC BY-NC-ND 4.0.

Please cite as: Lodge, J. M., Thompson, K., & Corrin, L. (2023). Mapping out a research agenda for generative artificial intelligence in tertiary education. *Australasian Journal of Educational Technology*, 39(1), 1-8. <https://doi.org/10.14742/ajet.8695>