GILE Journal of Skills Development

Generative AI Robs Students of the Joy of Learning

Ian Fellows

Canterbury Christ Church University, UK
ORCID: https://orcid.org/0000-0002-6754-3160

Abstract

Rapid developments in the field of Generative AI have caused businesses, educators and politicians to consider how best to accommodate and utilise these new technologies. This article explores the implications of using AI, particularly Large Language Models (LLMs), in the writing process. While accepting that artificial intelligence has many valid and valuable uses in other spheres of human existence, this article argues that using LLMs at any stage of the academic writing process beyond grammar-checking is detrimental to student learning. This article explains the importance of students mastering each stage of the Basic Model of Learning to Write Well (learning, thinking, writing), identifies higher-order thinking as the key objective of education and reminds readers of why learning is one of the most joyful activities a human can experience.

Keywords/key phrases: Generative AI, educational philosophy, learning, academic writing

1. Introduction

In business and education, Generative AI in the form of large language models (LLMs) such as ChatGPT, Llama 3, or Claude is considered an inevitability to be adopted and adapted to. Across the globe, nation-states are committing to projects that they hope will place their country at the forefront of these emerging technologies. The logic of adoption, that it is better to be on the front foot than compelled to catch up, is built on the presumption that AI will inevitably become ubiquitous. This race to uncritical adoption occurs despite many valid criticisms of Generative AI, particularly LLMs. Widely expressed concerns around LLMs relate to their proneness to inaccuracy (Shi *et al.,* 2025), hallucination (Li *et al.,* 2024*)*, threats to data security (Bellovin, 2024), inbuilt degeneration of training data (Shumailov, 2024), intellectual property theft (Kirchhübel & Brown, 2024), reproduction of racial and gender biases (An *et al.,* 2024) and, most important of all, disastrous environmental impact and contribution to the climate crisis (Ding & Shi, 2024). Without considering these concerns, I wish to add a significant one: Generative AI robs students of the joy of learning. LLMs respond to prompt questions by algorithmically constructing the most statistically likely answer based on their training data (Dixon, 2023). LLMs are, therefore, a markedly different category of software than enabling

technologies such as spelling and grammar checking since the purpose is not to clarify a user's original thoughts but to act as a proxy for original thought. This article contends that the joy of learning is to be found in the process of generating original thought.

2. Basic Model of Learning to Write Well

When I work with students in the initial stages of their academic development, I introduce them to my basic model of learning to write well (Figure 1 below).

FIGURE 1: BASIC MODEL OF LEARNING TO WRITE WELL (2025)
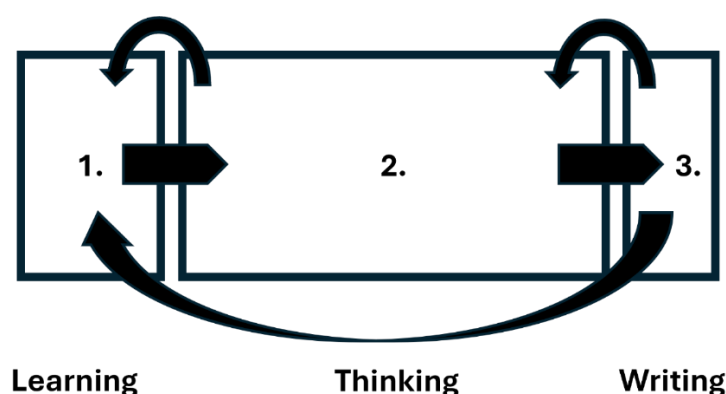
Source: Author's own, 2025

Box 1 represents Learning. This includes listening to lectures, asking questions, reading textbooks and journal articles, watching videos, doing simulations, and engaging in other information input activities. Box 2 is Thinking. As you can see, this is the largest box by a considerable distance. I attribute Box 1 20%, Box 2 70%, and Box 3 10%, not necessarily in terms of time but certainly energy and importance. Which leaves Box 3, Writing. Yes, in my model, only 10% of writing well results from writing. Writing, after all, is merely the expression of thought and thinking, Thinking, I tell my students sincerely, though perhaps pompously, is where the magic happens. There are arrows on the model which demonstrate that this is not a sequential and linear process: when thinking, we may realise there is more to learn; when writing, we may realise our thoughts are not clearly articulated or, perhaps, that there is a significant gap in our understanding. After all, as the great essayist Joan Didion once said, "I write entirely to find out what I'm thinking" (1976, p.2).

The Basic Model of Learning to Write Well is a framework that fosters critical thinking (Chatfield, 2022), deeper learning through meaning-making (Dolmans *et al*., 2016) and the making of associations (Ash *et al.,* 2012), eventually allowing students to assimilate these insights to existing schema or else accommodate knowledge in a new schema (Plant & Stanton, 2013). Generative AI interrupts this process in damaging ways. The most extreme instance is when students submit work wholly or generated mainly by an LLM, foregoing the three stages in my model altogether. However, this academic misconduct is universally recognised as illegitimate and not what I discuss here. Many argue that LLMs have legitimate academic uses, such as distilling complex information, planning or structuring writing, employing more persuasive language, and synthesising ideas. I contend that employing an LLM for these activities will damage students' engagement with knowledge and rob them of the joy of learning.

3. Using Generative AI in the Writing Process

Let us suppose that I ask students to write a critique of leadership articles written by industry experts. In moments, an LLM could provide a student with a summary, saving the student the 'effort' of reading the articles (or, let us be optimistic, of repeatedly re-reading them). This is one of the main tasks for which students could use Generative AI. To test this, I asked the most famous LLM ChatGPT to summarise the key themes emerging from my published journal articles (for ethical reasons, all my work is Open Access, so I was not engaging in nefarious copyright contravention for which ChatGPT has been heavily criticised). While broadly accurate, the resulting output was generic, superficial, and partial. For example, the theme of employability appears in all but one of those articles and even in the title of one, but ChatGPT somehow missed it. One could argue that I do not write well enough to get my point across, but that critique would need to be extended to my co-authors and the peer reviewers and editors involved in each publication. It is much more likely that the LLM output is not of an adequate standard.

This is not to say that artificial intelligence modelling does not have valid and valuable uses. Iterative AI can perform tasks that would otherwise be impossible or immoral, such as modelling the growth of cancerous cells without requiring the condition to go untreated (Beretta *et al*., 2024), predicting the maintenance requirements of machine tools (Lee *et al.,* 2019), enhancing agricultural efficiency by predicting environmental conditions (Linaza *et al.,* 2021), or optimizing drug delivery to combat infectious diseases (He *et al.*, 2021). Crucially, however, while these tasks can eliminate human weaknesses such as fatigue or distractedness, they are repetitive tasks that do not require cognition. But learning, thinking, and writing, if they are to be done well, do. Returning to our leadership assignment example, even if students wish to earnestly engage in the thinking stage, rather than asking an LLM to synthesise the articles, they will lack the context to critique the source materials effectively. Instead, let us consider a student critiquing Bennis and Thomas's seminal leadership article Crucibles of Leadership (2002). Bennis and Thomas argue that people become great leaders by enduring hardships and becoming stronger. Imagine our student reads the sentence, "It is the combination of hardiness and ability to grasp context that, above all, allows a person to not only survive an ordeal, but to learn from it, and to emerge stronger, more engaged, and more committed than ever" (2002, p.8) and is reminded of a documentary about Vietnam War veterans suffering post-traumatic stress disorder (PTSD), and remembers that the negative impacts of trauma may not manifest for several years and often present in unpredictable ways (Bonde *et al.*, 2022). Intrigued, our student returns to the learning phase to read about trauma responses and returns to reconsider the Bennis and Thomas article from a fresh, critical, inter-contextual perspective. Whatever the student writes now, whether it has the appropriate structure and grammatical correctness of an LLM output or not, will be immeasurably more original, more interesting and, quite simply, a better piece of academic work than one in which "the prose style of the output is both bland and uninspiring… characterised by being simplistic, having a lack of evidence… may contain simple errors… [and is] rather soulless" (Dixon, 2023, p.106).

4. Higher-Order Learning

It is worth considering what we hope to achieve in education and contrasting that with the purpose of LLMs. Bloom's taxonomy (Bloom *et al.,* 1956, p.122) dictates that developing higher-order learning means that "what we learn is intended for application to problem

situations in real life." Synthesising and applying knowledge inherently require the application of novel thought and the making of unpredictable connections. This is the opposite of what LLMs are programmed to do. ChatGPT will produce an algorithmically generated prediction of the most likely response to a prompt based on its training data. In other words, LLMs generate a predictable response that replicates what has gone before. Research that will be presented to the CHI Conference on Human Factors in Computing Systems in April 2025 states that "users with access to GenAI tools produce a less diverse set of outcomes for the same task, compared to those without. This tendency for convergence reflects a lack of personal, contextualised, critical and reflective judgement of AI output" (Lee *et al.*, 2025). Contrast this with the virtuous circle of learning, mastery, positive feedback, and enhanced intrinsic motivation when students apply themselves to a problem and create a solution they understand and can comprehensively articulate.

5. Testing Learning

Given the capacity for LLMs to respond to prompts, it is feasible to frame prompt questions that request responses that test a student's understanding of a topic they have learned. Ward *et al. (2024)* suggest that LLMs are an effective study aid and can improve academic performance when used in a complementary role, with testing tailored to the student's current level being a specific example. Prompting an LLM to produce Socratic questions (Paul & Elder, 2019) that guide students to find new ways to express thoughts already elucidated could also be beneficial to test learning. Crucially, in these examples, Generative AI will only effectively support learning that a student has engaged in fully. It is helpful to think of an LLM as a coach for students: an athlete will not become faster if the coach is doing all the work. Much like grammar checking, the intention should be to support and complement rather than interrupt or circumvent the learning, thinking, and writing process. However, given the environmental cost, unethical use of other people's creative output, and ongoing reliability issues with LLMs, this article does not endorse its usage.

6. Conclusion

This article has explored the deleterious impact of LLM usage on the academic writing process. While uses for LLMs that do not interfere with the learning process may be conceived, they do not add anything to the process that a peer, mentor, fellow student or tutor can already do. All things being considered, using LLMs at any stage of the academic writing process beyond grammar-checking or as a learning test cheapens education, reducing it to little more than a commercial transaction in which the resource-intensive elements are outsourced in the name of efficiency. The recommendation of this article and question for educators when considering the recommendation or use of Generative AI must surely be, 'Does this replace any part of the learning experience?' if the answer is yes, then it must be set aside. Education is a transformative act, a path to developing skills, attributes, and understanding, that unconstrained and ill-thought-out use of Generative AI threatens to short-circuit. Learning is one of the most emotionally nourishing activities we can experience as humans, and the only conclusion to be had based on all available evidence is that Generative AI robs students of the joy of learning.

References

An, J., Huang, D., Lin, C., & Tai, M. (2024). *Measuring Gender and Racial Biases in Large Language Models* (arXiv:2403.15281). arXiv. https://doi.org/10.48550/arXiv.2403.15281

Ash, I. K., Jee, B. D., & Wiley, J. (2012). Investigating insight as sudden learning. *The Journal of Problem Solving*, *4*(2). https://digitalcommons.odu.edu/psychology_fac_pubs/63/

Bellovin, S. M. (2024). DeGenerative AI? *IEEE Security & Privacy*, *22*(03), 88–88. https://doi.org/10.1109/MSEC.2024.3385549

Bennis, W. G., & Thomas, R. J. (2002). Crucibles of leadership. *Harvard Business Review*, *80*.

Beretta, E., Cavaterra, C., Fornoni, M., Lorenzo, G., & Rocca, E. (2024). *Iterative algorithms for the reconstruction of early states of prostate cancer growth* (arXiv:2409.12844). arXiv. https://doi.org/10.48550/arXiv.2409.12844

Bloom, B., Englehart, M. D., Furst, E. J., Hill, W. H., & Krathwohl, D. R. (1956). *Taxonomy of educational objectives: The classification of educational goals. Handbook I: Cognitive domain*. New York, NY: David McKay.

Bonde, J. P. E., Jensen, J. H., Smid, G. E., Flachs, E. M., Elklit, A., Mors, O., & Videbech, P. (2022). Time course of symptoms in posttraumatic stress disorder with delayed expression: A systematic review. *Acta Psychiatrica Scandinavica*, *145*(2), 116–131. https://doi.org/10.1111/acps.13372

Chatfield, T. (2022). Critical Thinking: Your Guide to Effective Argument, Successful Analysis and Independent Study (2nd Edition). London: Sage

Ding, Y., & Shi, T. (2024). Sustainable LLM Serving: Environmental Implications, Challenges, and Opportunities : Invited Paper. *2024 IEEE 15th International Green and Sustainable Computing Conference (IGSC)*, 37–38. https://doi.org/10.1109/IGSC64514.2024.00016

Dixon, S. (2023). Don't panic (yet): The implications of ChatGPT for Education Studies in the UK. *EducationalFutures*, *14*(1), 96–116. https://educationstudies.org.uk/?p=20864

Dolmans, D. H. J. M., Loyens, S. M. M., Marcq, H., & Gijbels, D. (2016). Deep and surface learning in problem-based learning: A review of the literature. *Advances in Health Sciences Education*, *21*(5), 1087–1112. https://doi.org/10.1007/s10459-015-9645-6

He, S., Leanse, L. G., & Feng, Y. (2021). Artificial intelligence and machine learning assisted drug delivery for effective treatment of infectious diseases. *Advanced Drug Delivery Reviews*, *178*, 113922. https://doi.org/10.1016/j.addr.2021.113922

Kirchhübel, C., & Brown, G. (2024). Intellectual property rights at the training, development and generation stages of Large Language Models. In I. Siegert & K. Choukri (Eds.), *Proceedings of the Workshop on Legal and Ethical Issues in Human Language Technologies @ LREC-COLING 2024* (pp. 13–18). ELRA and ICCL. https://aclanthology.org/2024.legal-1.3/

Lee, H.-P., Sarkar, A., Tanklevitch, L., Drosos, I., Rintel, S., Banks, R., & Wilson, N. (2025, May 26). *The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers*. in CHI Conference on Human Factors in Computing Systems (CHI '25), Yokohama, Japan. Retrieved February 10, 2025, from https://hankhplee.com/papers/genai_critical_thinking.pdf

Lee, W. J., Wu, H., Yun, H., Kim, H., Jun, M. B. G., & Sutherland, J. W. (2019). Predictive Maintenance of Machine Tool Systems Using Artificial Intelligence Techniques Applied to

Machine Condition Data. *Procedia CIRP*, *80*, 506–511. https://doi.org/10.1016/j.procir.2018.12.019

Li, J., Consul, S., Zhou, E., Wong, J., Farooqui, N., Ye, Y., Manohar, N., Wei, Z., Wu, T., Echols, B., Zhou, S., & Diamos, G. (2024). *Banishing LLM Hallucinations Requires Rethinking Generalization* (arXiv:2406.17642). arXiv. https://doi.org/10.48550/arXiv.2406.17642

Linaza, M. T., Posada, J., Bund, J., Eisert, P., Quartulli, M., Döllner, J., Pagani, A., G. Olaizola, I., Barriguinha, A., Moysiadis, T., & Lucat, L. (2021). Data-Driven Artificial Intelligence Applications for Sustainable Precision Agriculture. *Agronomy*, *11*(6), Article 6. https://doi.org/10.3390/agronomy11061227

Paul, R., & Elder, L. (2019). *The Thinker's Guide to Socratic Questioning: Based on critical thinking concepts and tools*. London: Rowman & Littlefield.

Plant, K. L., & Stanton, N. A. (2013). The explanatory power of Schema Theory: Theoretical foundations and future applications in Ergonomics. *Ergonomics*, *56*(1), 1–15. https://doi.org/10.1080/00140139.2012.736542

Shi, X., Liu, J., Liu, Y., Cheng, Q., & Lu, W. (2025). Know where to go: Make LLM a relevant, responsible, and trustworthy searchers. *Decision Support Systems*, *188*, 114354. https://doi.org/10.1016/j.dss.2024.114354

Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. *Nature*, *631*(8022), 755–759. https://doi.org/10.1038/s41586-024-07566-y

Ward, B., Bhati, D., Neha, F., & Guercio, A. (2024). *Analyzing the Impact of AI Tools on Student Study Habits and Academic Performance* (arXiv:2412.02166). arXiv. https://doi.org/10.48550/arXiv.2412.02166

Declaration Statements

Open Access Agreement

Corresponding Author

The corresponding author for this manuscript is Ian Fellows who can be contacted by email via ian.fellows@gile-edu.org.