


Generative AI and the unceasing acceleration of academic writing

*Despite the prospect and existence of AI generated texts having been around for some time, the launch of ChatGPT has galvanized a debate around how it could or should be used in research and teaching. Putting aside the ethical issues of using AI in academic writing, **Mark Carrigan** argues that the dynamic of ChatGPT and generative AIs as efficiency tools opens the door to further growth and acceleration in research outputs, but also raises questions about the value of these products of academic labour.*

It has been less than four months since OpenAI released ChatGPT, their chatbot built upon the GPT-3.5 large language model, but it already feels like we have been talking about generative AI forever. Their release was perfectly calibrated to generate a viral hit as users throughout the world shared screenshots of their often eerie, occasionally erroneous, conversations through social media. Its evident capacity to produce plausible answers to descriptive questions provoked anxiety throughout higher education, raising the spectre of familiar instruments of assessment being rendered redundant overnight.

To the extent scholarship has figured in these discussions, it has been restricted to the question of ChatGPT being cited as a co-author on papers. After at least four [instances](#) in which the chatbot was credited in this way, [Science and Springer Nature have prohibited the practice](#). This illustrates the speed with which norms surrounding the use of generative AI are being formed within the sector, as well as the path dependence which will govern their development over time. For example Science's editor-in-chief argued that text produced by Chat-GPT would be seen as plagiarism whereas Nature have allowed their use under specific documented circumstances. These early reactions are likely to establish the parameter in which future controversies will play out, with the release of GPT-4 and its integration into Microsoft Office likely to turbo charge these discussions.



While the formal attribution of authorship is clearly an important question, it raises a deeper issue about why authors would seek help from automated systems and what this suggests about the current state of academic publishing. When multiple generations of academics have internalised the imperative to [‘publish or perish’](#), how will they respond to a technology which promises to automate significant elements of this process? Is there a risk that the capacity to automate aspects of the writing process will simply lead to more writing? There are important issues remaining to be clarified about what constitutes acceptable uses of GPT within different domains of academic practice but there’s a broader challenge here in terms of how our conceptions of scholarly productivity have escalated in [an accelerating academy](#).

One [estimate from 2015](#) suggested around 34,550 journals published around 2.5 million articles per year. This [later study from 2018](#) found over 2.5 million outputs in Science & Engineering alone, highlighting how growth rates over a ten year period varied between 0.71% in the United States and 0.67% in the United Kingdom to 7.81% in China and 10.73% in India. Obviously there are factors at work here other than escalating expectations of scholarly output, such as the international growth of scientific fields and

the intellectual interconnections generated by the digitalisation of academic publishing. But, if we accept the premise that generative AI has the potential to automate parts of the writing process, then it increases how many outputs we can produce in the same amount of time. Imagine what annual outputs might look like globally if generative AI becomes a routine feature of scholarly publishing.

Why do we publish? In my experience academics can be weirdly inarticulate about this question. It is what we are expected to do and it is therefore what we do, often with little overarching sense of the specific goals being served by these outputs other than meeting the (diffuse or explicit) expectation of our employers. In a quantified academic world it is far too easy to slip into imagining countable publications as an end in themselves. These are conditions in which technologies which change the time:output ratio could prove extremely seductive. If this technology is taken up an individualised way, reflecting the immediate pressures which staff are subject to in increasingly anxiety ridden institutions, the consequences could be extremely negative. In contrast if we take this opportunity to reflect on *what* we might use this technology for as scholars and *why*, this could herald an exciting shift in how we work which reduces the time spent on routine tasks and contributes to a more creatively fulfilling life of the mind.

The content generated on this blog is for information purposes only. This Article gives the views and opinions of the authors and does not reflect the views and opinions of the Impact of Social Science blog (the blog), nor of the London School of Economics and Political Science. Please review our [comments policy](#) if you have any concerns on posting a comment below.

Image Credit: [DeepMind](#) via Unsplash.
