



# ‘A rather stupid but always available brainstorming partner’: Use and understanding of Generative AI by UK postgraduate researchers

Ross English , Rebecca Nash  and Heather Mackenzie 

Centre for Higher Education Practice, University of Southampton, UK

## ABSTRACT

Research into the increased use of Generative AI in Higher Education has largely focused on undergraduate study. While many institutions are grappling with the implications for doctoral level, there has been little published work investigating how postgraduate researchers use the technology or their attitudes towards it. This paper is based on a survey of 75 doctoral candidates across 19 UK Higher Education Institutions. The results show that most respondents had used Generative AI for their doctoral research, with the most common uses being framed as time-saver, editor or colleague. There was an awareness of limitations and ethical issues connected to the use of AI but no agreement as to where those boundaries lie. The paper concludes that there is an urgent need for sector agreement and communication on acceptable use and best practice.

## KEYWORDS

Generative AI; postgraduate research; doctoral study

## Introduction

The use of Generative AI in Higher Education has become an urgent issue since the launch of the freely available *Open AI* ChatGPT chatbot in November 2022. Generative AI are ‘deep-learning’ programmes that can generate human-like responses to user prompts (Lim et al., 2023). This includes the generation of images, computer code, or text. Generative AI is a significant advance on Conversational AI models. Rather than relying on pre-programmed responses to queries, Generative AI draws on information databases to create seemingly ‘new’ answers. While not the first Large Language Model (LLM), the arrival of ChatGPT-3 accelerated the debate around Generative AI and its potential for good or ill in education, as it provided a free service that was more accessible than most of its rivals. GPT-3 was different to previous versions, with the conversational style of the interface making the technology easy to navigate for the novice user. In addition to user-friendliness, it was far more powerful, with 175 parameters (learnable components that can be adjusted during the training process) and was trained on 45TB of text from a variety of sources, including Wikipedia (García-Peñalvo, 2023; Vázquez-Cano et al., 2023).

**CONTACT** Ross English  [r.m.english@southampton.ac.uk](mailto:r.m.english@southampton.ac.uk)  Centre for Higher Education Practice, University of Southampton, BJ, SO17 1 UK

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

Reaction from educators to the increased accessibility and rising popularity of Generative AI, particularly those that generate text responses, has been mixed. Common concerns have surrounded plagiarism in the form of students using LLMs to write assessed work, and the (in)ability of staff or plagiarism software to identify it. Some studies have suggested that not only Generative AI-written assignments can avoid detection by staff but can also be of sufficient quality to earn high grades (Bommarito & Katz, 2022; Katz et al., 2024; Vázquez-Cano et al., 2023). However, elsewhere, researchers have found AI-generated work to be either inaccurate (Gregorcic & Pendrill, 2023) or easily detected by existing plagiarism software (Aydin & Karaarslan, 2022; Gao et al., 2023). The reaction to this perceived threat to academic integrity led to a diverse range of approaches and policies across Higher Education, ranging from prohibition and discouragement to, increasingly, cautiously embracing the technology's potential (McDonald et al., 2024).

The initial flurry of publications on this topic has been largely undergraduate or taught postgraduate-focused, centring on the question of how Generative AI impacts formal assessment. Published work on the impact of doctoral-level study has largely been absent (an exception is Wang et al., 2018, which predates ChatGPT and its competitors). This is not to say that work on taught degrees is irrelevant to those involved in postgraduate research provision. Indeed, many doctoral students are required to submit written work for assessment at key points in their journey. One integral aspect missing from the literature is the voice of the postgraduate researcher, both in terms of their attitudes towards the use of the technology in doctoral study and evidence as to how it is already being used. Not surprisingly, studies which have examined attitudes of taught students towards Generative AI found a range of (dis) comfort about including Generative AI-produced writing in essays and the extent to which such work could be claimed as students' own (Fyfe, 2023). Tlili et al. (2023) found students aware of the limitations and potential, but some were quite open about their reliance on AI when answers to questions were not readily available. It is not clear whether such findings reflect mixed attitudes towards plagiarism generally (Kattimani et al., 2022) or are explicitly linked to an evolving relationship with the new technology.

Debates around Generative AI in academic research and publications, and therefore relevant to doctoral degrees, echo many of the themes present in the literature on undergraduate study: the benefits offered versus the threat of plagiarism. For researchers who work with large datasets, complex models, simulations, or coding, the efficiencies offered by existing or future AI technology may be an enormous boon (Alshater, 2022). There is more uncertainty regarding where and when LLMs can acceptably assist with writing academic research. Where are the lines to be drawn between using Generative AI to improve the standard of a researcher's writing, generating an abstract from an already written article, contributing to body text or even being credited as a co-author (Dwivedi et al., 2023; Tlili et al., 2023)? The fundamental difference between the questions surrounding the use of Generative AI by taught students and those concerning researchers, be they staff or postgraduate, is the issue of originality. LLMs can create an apparently 'new' piece of writing but, in reality, are simply predicting the most appropriate sequence of words from existing text in their database, adjusted based on interactions with the user (García-Peñalvo, 2023). It cannot create original research, either theoretical or empirical –

the essence of a doctorate, and something which doctoral researchers are expected to demonstrate.

With doctoral study occupying a distinct space within Higher Education, encompassing both the world of study and of research, particular attention needs to be given to this cohort when considering policy or practice in any area, including Generative AI. We argue that without seeking insight into the views and practice of doctoral researchers before considering their own approach, institutions are in danger of shaping policies in the dark, advising or regulating in line with leaders' own preconceptions and not regarding current realities. This study therefore aimed to describe why and how current postgraduate researchers in the UK are using Generative AI in their doctoral research and writing.

## Methods

Postgraduate Researchers (PGRs; often also referred to as doctoral students or candidates) from across the UK were invited to complete an anonymous online questionnaire hosted on the Qualtrics XM survey platform. Potential participants were invited to view an information sheet and consent form via the survey platform and to complete the questionnaire only if they were happy to participate. The Participant Information Sheet, consent section, and final section of the survey made it clear that consent could be withdrawn up to the point of final submission but that, once submitted, the anonymous nature of the survey meant that recorded responses could no longer be removed. Participants were reminded that they could contact their Graduate School/Doctoral College to find out about guidance/training on academic integrity if, they had concerns about their own use of Generative AI.

The questionnaire comprised five demographic questions and 21 questions about awareness and use of Generative AI and took approximately 20 min to complete (see <https://doi.org/10.5258/SOTON/D3304>). The survey was promoted via two national organisations supporting PGRs and via informal contact with Graduate Schools and Doctoral Colleges. The survey was additionally promoted via the authors' social networks on X and LinkedIn. It opened on 4 December 2023 and closed on 28 February 2024. The quantitative data was analysed using IBM SPSS Statistics 29.0, and the qualitative data from the free-text questions were analysed using reflexive thematic analysis (Braun & Clarke, 2006) using NVivo 14. The qualitative results were analysed separately by two of the authors and then compared and refined during subsequent conversations. Analysis was largely inductive due to the exploratory nature of the survey. No Generative AI tools were used in the research or writing of this paper.

## Results

### *Sample description*

In total, 75 PGRs completed the survey from 19 UK HEIs (four in London, four in the South-East, two in Yorkshire and the Humber, two in the South-West, two in the East Midlands and one in each of the West Midlands, Wales, North-West, Northern Ireland and Scotland) of which seven were post-1992 universities. PGRs were studying in a range of disciplines.

**Table 1.** Participant characteristics.

| Characteristics                 | Number (%) |
|---------------------------------|------------|
| <b>Age</b>                      |            |
| 18–24                           | 7 (9.3)    |
| 25–34                           | 38 (50.7)  |
| 35–44                           | 20 (26.7)  |
| 45–54                           | 8 (10.7)   |
| 55 and over                     | 2 (2.7)    |
| <b>Gender</b>                   |            |
| Female                          | 46 (61.3)  |
| Male                            | 28 (37.3)  |
| Non-binary                      | 1 (1.3)    |
| <b>English primary language</b> | 52 (69.3)  |
| <b>Discipline studying</b>      |            |
| Social Sciences                 | 33 (44.0)  |
| Physical & Life Sciences        | 12 (16.0)  |
| Engineering & Mathematics       | 10 (13.3)  |
| Arts & Humanities               | 9 (12.0)   |
| Business & Law                  | 7 (9.3)    |
| Medicine                        | 4 (5.3)    |

**Table 2.** Whole sample use and understanding of Generative AI.

| Question  | Yes: n (%)   | No: n (%) | Not sure: n (%) |
|---|--|-----------|-----------------|
| Heard of Generative AI  | 73 (97.3)  | 2 (2.7)   | N/A             |
| Used Generative AI for any purpose  | 57 (76.0)  | 18 (24.0) | N/A             |
| Interested in the potential of Generative AI for academic work                                | 47 (62.7)  | 12 (16.0) | 16 (21.3)       |
| Used Generative AI in any aspect of their doctoral research/writing                           | 39 (52.0)  | 18 (24.0) | 18 (24.0)       |
| Received training or been supplied with information on local policy by own HEI                | 24 (32.0)<br>(Training only = 2 (2.7); Information only = 20 (26.7); Both = 2 (2.7)) | 51 (68.0) | N/A             |
| Believe understand data security issues related to Generative AI in doctoral research/writing | 38 (50.7)  | 16 (21.3) | 21 (28.0)       |

**Table 3.** Further description of Generative AI use among those who had used within their doctoral research/writing ( $n = 39$ ).

| Question   | Yes: n (%)                                  | No: n (%)                                  | Not sure: n (%)                       |
|--|---|--|---------------------------------------|
| Included text generated by AI in academic writing related to their doctoral research | 9 (23.1 which is 12.0 of the whole sample*) | 30 (76.9 which is 40 of the whole sample*) | 0 (0.0)                               |
| Included text in work submitted for assessment related to doctoral degree            | 7 (17.9 which is 9.3 of whole sample*)      | 2 (5.1 which is 2.7 of whole sample*)      | 0 (0.0)                               |
| Made it clear to reader that Generative AI tools had been used in relation to above  | 0 (0.0)                                     | 6 (15.4 which is 8.0 of whole sample*)     | 1 (2.6 which is 1.3 of whole sample*) |
| Included text in work submitted for public consumption                               | 2 (5.1 which is 2.7 of whole sample*)       | 7 (17.9 which is 9.3 of whole sample*)     | 0 (0.0)                               |
| Made it clear to reader that Generative AI tools had been used in relation to above  | 0 (0.0)                                     | 1 (2.6 which is 1.3 of whole sample*)      | 1 (2.6 which is 1.3 of whole sample*) |

\*Of whole sample ( $N = 75$ ).

The demographic characteristics of respondents were skewed towards females (61.3%) and the Social Sciences (44.0%) (see [Table 1](#)).

### *Use and understanding of Generative AI*

Of the sample, 97.3% ( $n = 73$ ) had heard of Generative AI, and 76.0% ( $n = 57$ ) had used it (for any purpose). About 62.7% ( $n = 47$ ) reported interest in the potential of Generative AI for academic work, 16.0% ( $n = 12$ ) said they were not interested and 21.3% ( $n = 16$ ) said they were not sure (see [Table 2](#)).

Thirty-nine (52.0%) PGRs reported having used Generative AI within their doctoral research/writing. Of those, 23.1% had included AI generated text within academic writing related to their doctoral research (12.0% of the whole sample) and 17.9% had included this text in work submitted for assessment related to their doctorate, including progression reviews or equivalent (9.3% of the whole sample). None of the seven PGRs who included such text in the work submitted for doctoral assessment reported that they made it clear to the reader that Generative AI tools had been used (six reported they did not, one answered 'not sure') (See [Table 3](#)).

Of the 36 PGRs who reported that they had either not used Generative AI within their doctoral research/writing or were not sure, only four (11.1%) said that they would ever consider doing so with 20 (55.5%) saying they would not and 12 (33.3%) being unsure.

ChatGPT was the tool most heard of by participants and the most used tool for both any purpose and specifically for doctoral research and/or writing, where it was the dominant tool in use (by 46.7% of the whole sample, compared with 4.0% who had used Bard (now Gemini), the next most frequently used Generative AI tool within doctoral research and/or writing; no other tool registered more than one user). In total, participants reported having used at least one of 23 Generative AI tools for any purpose, and one of the 20 in their doctoral research and/or writing.

## **Free text responses**

### *PGR perceptions of Generative AI*

Attitudes amongst the respondents towards Generative AI ranged from functional to personal. There were diverse opinions about the technology's reliability. The most used word to describe the technology was 'tool' (46 times in free-text responses) with many PGRs viewing Generative AI as simply a technological resource like any other:

This is a tool, the same as Word, Excel, Google, WhatsApp, YouTube and others.

It's a tool – like Word or Excel. If trained properly, it can be used for any purpose.

A common comparison was with Google, which fitted a view of LLMs as just another gateway to internet searches, and as such, nothing radically different from what they used before:

Before Google people had to look in a library and manually use an index. Search engines have made gathering information easier. Generative AI can do the same.

Framing Generative AI in evolutionary terms against previous – more cumbersome and time-consuming – methods of information-seeking gives an impression of inevitability. In reference to Google, there was emphasis on Generative AI as an improvement on the search engine:

Instead of going to Google and spending a long time searching for the info myself, I asked ChatGPT.

The consequence of this was that, viewed as a tool like those that have become an everyday part of a researcher's life, there was often a sense of inevitability that Generative AI would quickly become the same. This led to conclusions that soon the ability to use this technology, whether in research or other professions, would become a skill that is necessary to develop. In this way, PGRs would need to be encouraged to learn how to incorporate Generative AI platforms as part of their developing skill set:

All PhD students should be supported in their engagement with AI just as they would be supported with other tools such as use of data analysis software, social media and online/cloud file management.

A few participants also raised the possibility of Generative AI as something that could help level the playing-field for PGRs who would otherwise not have the advantages available to others:

I think this technology has a range of huge applications in research, specifically for disabled researchers, and researchers who work in institutions with less administrative support than others.

However, for all the perceived usefulness of the technology, there was, for some, an awareness of its limitations, particularly around reliability:

These tools are normally very unreliable, and not designed for research.

At the moment it is inaccurate and unreliable.

This suspicion sometimes led to wholesale rejection of the technology:

I simply don't support it - I can't see the current generation of generative AI software as anything other than technology for technology's own sake ... I believe generative AI feeds into a "productivity culture" where everything in life is task- and progress-oriented and is actually dehumanizing us/chaining us down rather than freeing us to expand our lives.

Not sure there is much of a potential, other than having a rather stupid but always available brainstorming partner.

This last quote introduced another emerging theme: while many viewed LLMs as analogous to existing online tools, for others, the chat-bot interface changed how they viewed their relationship with the technology. There was a feeling of AI either taking on human characteristics or fulfilling roles normally played by people:

AI has real potential to be a great colleague.

[It is] easier than asking [my] supervisor for support.

This seemed important, as the distinction between a tool and a colleague is that the former is used by the researcher, whereas the latter is an active participant and considered so by the user. Attention will now be turned to explore broader aspects of the research process that PGRs used Generative AI for, building on the perceived benefits and drawbacks.

### *Efficiency, clarity and problem-solving: Generative AI and research tasks*

Building on the usefulness of Generative AI as research tool(s), respondents pointed to enhanced efficiency as a key benefit across multiple aspects of the research process, particularly earlier on regarding literature. Searching for relevant literature was mentioned as an advantage of using Generative AI:

To cut down on time sifting through literature, it helps me narrow my searches for relevant information.

Looking at currently researched topics it makes research more accessible because there are so many publications and so much work being produced often searching papers is overwhelming.

Cutting out ‘overwhelming’ elements of early research, i.e. where to start and how to find the most relevant and/or up-to-date literature on a topic, was felt to reduce the time that may otherwise be spent exploring and potentially disregarding publications of little use. There was awareness, however, that the output of some Generative AI tools regarding references may not be ‘real’:

I experimented with asking chat GPT about what follow-up papers had been written on a first paper. I realised very quickly it made papers up.

In echoing the previous limitation outlined by PGRs regarding unreliability, acknowledging that ChatGPT outputs, in particular, may comprise fabricated references removes some of the ease with which ‘useful’ literature can be sought out quickly. Efficiency in bypassing search engine literature results is potentially countered by the time taken to substantiate the existence and/or accuracy of Generative AI output. As another respondent highlighted:

I find ChatGPT helpful for quick assistance, though I always crosscheck the information with other sources. Before relying on ChatGPT, I used to search across multiple platforms, which consumed a lot of time.

It is not clear how much time is saved by the individual when they are having to double-check ChatGPT output to ensure reliability, or whether this considerably outweighs using another tool for the job.

Going beyond locating literature, Generative AI was used to provide summaries of text that PGRs may be unfamiliar with, or puzzled by, to clarify their understanding:

I have used generative AI to summarise complex bits of text (paragraphs not full chapters) on academic papers I am reading if the contents of said paper are particularly dense.

I used ChatGPT to interpret a paragraph of text for me as I wanted to confirm my understanding of what the text was saying.

Using Generative AI to elucidate meaning from difficult text acts as a ‘second opinion’, emphasising the use of Generative AI as a colleague or collaborator. There was, however, some elaboration on boundaries of acceptable use when using Generative AI to clarify information:

... if we don’t understand a subject/theory/idea, it is acceptable to use ChatGPT to give us a baseline understanding, although we are then expected to expand on that knowledge ourselves using appropriate methods such as reading journal articles.

Generative AI could be used akin to a dictionary, glossary or textbook to aid with understanding, which sets the individual up effectively to pursue tasks like reading literature with a greater knowledge of ‘baseline’ information.

The other frequently cited use of Generative AI in research tasks was its use to generate programming code. Respondents cited using Generative AI to help with R programming language. Again, this was framed as time-saving:

It speeds up programming (proposes new lines, which saves time on typing, can write a piece of generic code from a simple prompt). It can also sometimes solve errors.

I have also used ChatGPT to help me write code (R and JavaScript). e.g. I wanted to create a nice-looking output table for some descriptive statistics and I didn’t know how to do this.

Simply reducing the time it takes to type out code is a benefit, particularly if the prompt is short. In producing visually pleasing outputs that may be difficult or time-consuming to realise, Generative AI was felt to lessen learning time and speed up outputs. This further extended to Generative AI fixing code:

[Fixing] my coding errors in a way that maximises amount of time that I can spend on my research rather than trying to fix coding bugs.

Maximising time spent on other research tasks interestingly separates out what elements of research are felt to be ‘worthy’ of expending time and effort on, and those that are perhaps an ‘important nuisance’ (i.e. a broken piece of code), and a barrier to research progression.

Generative AI was felt to be both a help and a hindrance. Where it could be useful (finding text, providing clarity for complex text, fixing code) was countered by recognised limitations (inaccurate outputs, made up references), which could indeed remove some of the perceived efficiencies. Looking to doctoral writing, a similar picture emerged: one of the perceived usefulness for functional language checks and ‘improving’ writing for an academic audience. However, concerns abound not only regarding stylistics of output but also boundaries of acceptable use when considering the pursuit of originality and where it is deemed important for PGRs to focus their academic efforts.

### ***A trusted editor? Generative AI and doctoral writing***

Generative AI was perceived as a useful editor for a PGR’s work, particularly where English may not be their first language. It was used less frequently to generate text beyond basic planning. Respondents who used Generative AI for writing felt that it was helpful for

‘improving’ their writing, particularly regarding conciseness, grammar, and achieving an ‘academic tone’:

I use ChatGPT to double check my writing. I have an open tab where I have trained GPT to understand my style of writing. There I ask ChatGPT to read paragraphs or sections of my work, and give me an improved version that is clearer, more engaging, yet academic.

I have submitted previously written paragraphs to GPT with a set of specific instructions to proof-edit the text in order to shorten sentences and/or make them more readable based on the specified criteria.

Here, again, are allusions to efficiency – some respondents ‘training’ ChatGPT to understand their writing style or using ‘specific’ prompts to generate a desired output, whereas others would take output and rewrite it to ensure it maintained their linguistic ‘style’. Some respondents were more critical of Generative AI’s ability to produce anything that was stylistically distinctive (or functional) and went beyond surface-level:

... the current ChatGPT is very waffle-y so I think it’s quite obvious AI is being used and makes everyone’s work sound the same.

I’m less interested in the plagiarism arguments around this stuff than a) the fact that everything it outputs is utter shite (meaningless drivel, structurally flimsy, stylistically awful).

Critics of Generative AI’s value in the writing process noted limitations in both copy-editing and producing texts that exceeded the academic competencies of the user who generated the prompt. Here, though, it should be noted that there are likely differences between those who are using Generative AI to check over their functional language skills, and those who are hoping that the tools may be able to create an engaging text for a specific purpose. Regarding the latter, it was found to be less useful than hoped:

I have tried using both Notion and ChatGPT to give me an abstract for the title of a talk or poster I was doing when I was struggling to bring my ideas together. But I didn’t find it useful [...] the AI just gave me a very basic and normative abstract.

Beyond the functional and stylistic limitations of Generative AI for doctoral writing, there was unease about the use of Generative AI in the writing process:

In academic writing, using ChatGPT for English proofreading can be beneficial for catching grammatical errors and improving overall writing clarity. However, relying solely on it to generate entire paragraphs or sections for research writing would indeed be considered unethical, as it could compromise the integrity of the work and violate academic standards.

Boundaries for acceptable use were put in place by respondents who felt that using Generative AI as a more intuitive spelling and grammar checker was largely fine, but that prompting the tools to create text relevant to the thesis was beyond the scope of unproblematic usage:

I think research is mostly fine and limited writing i.e. changing a title to sound more interesting or writing a more relevant abstract is okay using it to improve work seems okay but I think there should be limitations on using it getting it to write parts of a publication/thesis feels wrong because it feels like it’s not my work.

Much like use of Generative AI in research support, using such tools in the process of doctoral writing uncovered a spectrum of what was considered appropriate boundaries of use. There seemed to be a consensus that Generative AI should not be used to create large amounts of text for the author, as that would denigrate their credibility as a developing scholar. However, the use of Generative AI for functional language checks and to make adjustments to pre-written work with the aim to 'improve' it for an academic audience was deemed more acceptable.

### *Ethical concerns and need for regulation*

Following on from what PGRs felt to be appropriate boundaries of use, some raised allied concerns regarding the ethics of using ChatGPT in research. Ethical concerns fell into three categories: plagiarism, data security, and algorithmic-bias. There were some references to the potential for the text generated by LLMs to either be misused as direct plagiarism or indirectly due to the uncertainty of the provenance of the generated text:

Because AI is trained on the works of others with no citations which in a way could be considered a form of plagiarism except even the author does not know what the original source is.

I've never been able to view Gen-AI as something that magically "produces" content, but instead as a commodified theft of other people's content.

However, the most common ethical concern mentioned was in the opposite direction: the worry of LLMs retaining researchers' inputted material and using it to reply to other users' queries. These concerns included issues of both intellectual property and data-security:

Since Generative AI is a deep-learning model, which produces answers/solutions based on the big data that it has collected (also I assume that the AI would not yet be able to filter information such as personal data), we cannot certainly say that using AI for doctoral research is free from a risk of personal data breaches.

My concern would be that my original research findings then become part of the general information online. This may impact on my ability to publish.

A third concern was that the design of the algorithm that governed LLMs response to queries could in-bed biases into responses that may then be incorporated into people's research:

The second main issue relates to structural inequalities, and the increasingly known issues of AI perpetuating stereotypes of (e.g.) whiteness, cis-genderedness, ableism, sexism etc through uncritical learning and reproduction of unequal or biased content.

How might AI continue to perpetuate "whiteness" and a westernised canon and approach to "academic writing"? How might it stymie more creative and disruptive approaches? Will those with "non-normative" englishes or communication styles feel pressured into using AI to "correct" their writing, and what loss does that mean for diverse expression and voices within the academy.

For many respondents, these ethical concerns were threats rather than inevitable consequences of the technology. Where this position was taken, the predominant view was

that whether the use of Generative AI was ethically problematic or not was in the hands of the researcher themselves. 'It depends' was a commonly used phrase:

I think it depends what you're using it for. If you're asking it to write something from scratch, which you would then use in your work, then it's not your own writing/creation, so I find that completely unethical. In terms of editing, I believe it is acceptable because I am inputting my own text.

Depends how they are using AI. If it's for writing, then that's problematic. However if it's a support tool then cannot see this as an issue (although cannot be relied upon).

The conclusion reached by some was that clear guidance was needed quickly:

I am not aware if the university has a policy on the use of AI in doctoral research/writing (or indeed at any other levels of study at the institution). If it has, it needs to be publicised more widely, and if it hasn't got one, it needs to create one pronto!

## Discussion

This survey sought to establish how postgraduate researchers in UK HEIs use Generative AI and what viewpoints exist regarding its possibilities and future use. From the sample, it was clear that a majority had used Generative AI either generally or specifically for their doctoral research and/or writing. One caveat must be acknowledged, however: while the survey was promoted to PGRs regardless of whether they were current users of Generative AI, it could be the case that those who are already familiar with or are existing users may have been more likely to take the survey. Nevertheless, the responses revealed some interesting patterns. The most common form of Generative AI used in doctoral research was the text-generating LLMs, with Chat-GPT being by far the most utilised. Respondents reported a variety of uses. The three most acknowledged were as time-saver, editor, and colleague.

The concept of Generative AI as a labour-saving device raises an interesting issue: tacit acknowledgement that certain tasks within research are, for some, considered mundane or sufficiently unimportant to the process of research that they can be outsourced to a third-party. These tasks include code generation, table creation, literature searches, information summarising, and abstract writing. It was not clear from the responses whether there was any common agreement as to where the limits of this lies. A similar issue concerns the use of LLMs as text-editors. Several respondents were open about using Generative AI to re-write text they had produced, to proof-read, re-write, or suggest new material such as titles or abstracts. Others shared concerns about such use both in terms of the ability of LLMs to produce the necessary quality of writing, and the ethical issues raised. There seemed little agreement regarding which point using AI to contribute to written output crossed a line into plagiarism. The other ethical concern regards data security, particularly what could happen to original research that was inputted into an LLM. Of interest here was the dominance of Chat-GPT as the LLM of choice. Many HEIs recommend the use of Microsoft Bing Co-pilot as any information inputted into the chatbot can be ringfenced to within that institution as part of their Microsoft subscription,

a safety feature not available in ChatGPT. Despite this, Bing came low down in reported use.

The perception of AI as a colleague or substitute supervisor that can be used to help generate ideas and as a source of support or reassurance was more surprising. This raised some interesting questions. Because LLMs cannot produce new research, only rearrange existing information stored in its database, there runs a danger of ideas based on outputs from AI being perceived as new, but without appropriate referencing. Similarly, one of the key roles of research supervision or collegiality is to encourage the creation of new thought. There are also concerns around the provision of academic and pastoral support for PGRs if LLMs are being turned to as a substitute. However, it was not clear from the responses if those who used LLM as colleagues or supervisors were doing so in addition to or in place of existing support.

Twelve per cent of respondents confirmed that they had included text generated by AI in academic writing related to their doctoral research, but none of those had made that fact clear to the reader. While that is of concern, it raised a wider issue where there may be insufficient clarity: at what point or with what type of contribution would it become incumbent on a researcher to acknowledge the input of Generative AI? Is the inclusion of AI-generated text the appropriate boundary? It is now usual for academic publishers to include Generative AI use in their policies for authors, commonly but not exclusively, allowing limited use, declared to the reader, but not attributed as co-author (e.g. Taylor & Francis (n.d.)). As AI use becomes widespread, acknowledgement should become a standard part of methods sections, and it may be appropriate to begin to ask for such clarification in all work, including when AI has not been used, as we did in the methods section above.

As several respondents noted, the appropriate use of Generative AI may be a skill that will be needed to be developed by future researchers. As prohibition of the use of AI in research is not realistic or probably desirable, there is a need for clarity in guidance, regulation, and training. This appears not to be the case currently, as 68.0% of respondents reported that they had not received training or information on appropriate use. We acknowledge, however, that in a quickly evolving field, HEIs are providing more information as they, too, become more familiar by the capabilities and limitations of Generative AI.

Based on the responses received, the following areas seem to be surrounded by sufficient uncertainty to require clarification by HEIs and by the sector as a whole:

- Is the inclusion of AI-generated text, code, and images in work submitted for assessment or publication ever acceptable? If it is, when would it be considered acceptable?
- Are current publisher or institutional policies on whether, how, or in what circumstances, Generative AI be given recognition or authorship fit for purpose and whether sector-wide agreement is needed?
- Should the use of LLM as a proof-reader be allowed for work submitted for assessment or publication?
- What are the data-security issues surrounding the entry of research writing into Generative AI? How should considerations of data-security issues be taken account

of in research design, including within applications for ethical review and data management plans?

- Should, and if so in what circumstances, PGRs use unattributed information gathered from the use of AI?

What is noticeable about several of these issues is that they seem strikingly familiar. Were 'AI' replaced by 'supervisor' or 'colleague', a case could be made that there is a similar lack of clarity across the Higher Education sector. For instance, there is little consistency of approach in when and whether outside proof-readers can be used for doctoral theses and, in many institutions, little internal clarity (Harwood, 2019). At what point does supervisory support cross into co-authorship (Hockey, 1994)? Indeed, where are the boundaries between single-authored original work (the essential criteria for successful doctoral study) and collaboration (Katz & Martin, 1997)? In short, what is original research? While the appropriate use of Generative AI needs clarity, it should be used as an opportunity to clarify other, related, issues concerning doctoral study.

## Conclusion

This survey into the use of Generative AI by PGRs suggests that the technology is already widely used in doctoral study but that there is little clarity as to what the appropriate limits are. As the debate about Generative AI use across different levels of higher education continues apace, there are a mixture of positions. Either institutions should educate students as to the acceptable limits of use, adjust assessment methods and be alert for the tell-tale signs of AI generated text (Cotton et al., 2023) or to move away entirely from a 'police-catch-punish approach' (Kramm & McKenna, 2023) to one where it is embraced as a potentially positive transformative resource (Lim et al., 2023). The need for sector agreement on such issues, translated into guidance and training for researchers, is imperative. To do that, wider issues around collaboration and originality need to be similarly clarified.

## Acknowledgments

The authors would like to thank all who completed the survey and colleagues at institutions who helped promote the survey.

## Disclosure statement

No potential conflict of interest was reported by the author(s).

## Notes on contributors

*Ross English* is Senior Teaching Fellow in Academic Practice in the Centre for Higher Education Practice at the University of Southampton. He works with the Doctoral College, delivering training to postgraduate researchers. His research focuses on the experience of postgraduate researchers in UK Higher Education.

**Rebecca Nash** is Senior Teaching Fellow in Academic Practice in the Centre for Higher Education Practice at the University of Southampton. She has worked in academic development since 2016, having previously focused on academic research skill enhancement of students prior to entering Higher Education, and is currently with postgraduate researchers.

**Heather Mackenzie** is Doctoral College Professional Development Programme Manager in the Centre for Higher Education Practice at the University of Southampton. Heather has been working in postgraduate researcher in professional development for over a decade. In this time, she has led professional development programmes, facilitated a wide range of workshops for postgraduate researchers and acted as a doctoral supervisor.

## ORCID

Ross English  <http://orcid.org/0000-0001-9056-8158>

Rebecca Nash  <http://orcid.org/0009-0000-1819-1114>

Heather Mackenzie  <http://orcid.org/0000-0002-5241-0007>

## Ethical approval

Ethical approval was granted by the Faculty of Social Sciences Ethics Committee, University of Southampton. ERGO number: 89558

## References

- Alshater, M. M. (2022). Exploring the role of artificial intelligence in enhancing academic performance: A case study of ChatGPT. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4312358>
- Aydin, Ö., & Karaarslan, E. (2022). OpenAI ChatGPT generated literature review: Digital twin in healthcare. In Ö. Aydin (Ed.), *Emerging computer technologies 2* (pp. 22–31). Izmir Akademi Dernegi. <https://doi.org/10.2139/ssrn.4308687>
- Bommarito, M., II, & Katz, D. M. (2022). GPT takes the bar exam. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.4314839>
- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Cotton, D. R. E., Cotton, P. A., & Shipway, J. R. (2023). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International*, 61(2), 228–239. <https://doi.org/10.1080/14703297.2023.2190148>
- Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A. S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L. & Wright, R. (2023). Opinion paper: ‘So what if ChatGPT wrote it?’ multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642. <https://doi.org/10.1016/j.ijinfomgt.2023.102642>
- Fyfe, P. (2023). How to cheat on your final paper: Assigning AI for student writing. *AI & Society*, 38(4), 1395–1405. <https://doi.org/10.1007/s00146-022-01397-z>
- Gao, C. A., Howard, F. M., Markov, N. S., Dyer, E. C., Ramesh, S., Luo, Y., & Pearson, A. T. (2023). Comparing scientific abstracts generated by ChatGPT to original abstracts using an artificial intelligence output detector, plagiarism detector, and blinded human reviewers. *NPJ Digital Medicine*, 6(1). <https://doi.org/10.1038/s41746-023-00819-6>
- García-Peñalvo, F. J. (2023). The perception of Artificial Intelligence in educational contexts after the launch of ChatGPT: Disruption or Panic? *Education in the Knowledge Society*, 24, e31279. <https://doi.org/10.14201/eks.31279>

- Gregorcic, B., & Pendrill, A.-M. (2023). ChatGPT and the frustrated Socrates. *Physics Education*, 58(3), 035021. <https://doi.org/10.1088/1361-6552/acc299>
- Harwood, N. (2019). 'I have to hold myself back from getting into all that': Investigating ethical issues associated with the proofreading of student writing. *Journal of Academic Ethics*, 17(1), 17–49. <https://doi.org/10.1007/s10805-018-9322-5>
- Hockey, J. (1994). Establishing boundaries: Problems and solutions in managing the PhD supervisor's role. *Cambridge Journal of Education*, 24(2), 293–305. <https://doi.org/10.1080/0305764940240211>
- Kattimani, S., Kujur, A., Nagarajan, P., & Zayabalaradjane, Z. (2022). Attitude toward plagiarism among postgraduate students of a medical institute in South India. *Journal of Psychiatry Spectrum*, 1(2), 106–110. [https://doi.org/10.4103/jopsys.jopsys\\_35\\_22](https://doi.org/10.4103/jopsys.jopsys_35_22)
- Katz, D. M., Bommarito, M. J., Gao, S., & Arredondo, P. (2024). Gpt-4 passes the bar exam. *Philosophical Transactions of the Royal Society A*, 382(2270), 20230254. <https://doi.org/10.1098/rsta.2023.0254>
- Katz, J & Martin, B R. (1997). What is research collaboration?. *Research Policy*, 26(1), 1–18. [https://doi.org/10.1016/S0048-7333\(96\)00917-1](https://doi.org/10.1016/S0048-7333(96)00917-1)
- Kramm, N., & McKenna, S. (2023). AI amplifies the tough question: What is higher education really for? *Teaching in Higher Education*, 28(8), 2173–2178. <https://doi.org/10.1080/13562517.2023.2263839>
- Lim, W. M., Gunasekara, A., Pallant, J. L., Pallant, J. I., & Pechenkina, E. (2023). Generative AI and the future of education: Ragnarök or reformation? A paradoxical perspective from management educators. *International Journal of Management Education*, 21(2), 100790. <https://doi.org/10.1016/j.ijme.2023.100790>
- McDonald, N., Johri, A., Ali, A., & Hingle, A. (2024). Generative artificial intelligence in higher education: Evidence from an analysis of institutional policies and guidelines. *arXiv Preprint arXiv: 2402.01659*. <https://doi.org/10.48550/arXiv.2402.01659>
- Taylor & Francis. (n.d.). 'AI policy'. Taylor & Francis. <https://taylorandfrancis.com/our-policies/ai-policy/>
- Tlili, A., Shehata, B., Adarkwah, M. A., Bozkurt, A., Hickey, T. T., Huang, R., & Agyemang, B. (2023). What if the devil is my guardian angel: ChatGPT as a case study of using chatbots in education. *Smart Learning Environments*, 10(1), 15. <https://doi.org/10.1186/s40561-023-00237-x>
- Vázquez-Cano, E., Ramírez-Hurtado, J. M., Sáez-López, J. M., & López-Meneses, E. (2023). ChatGPT: The brightest student in the class. *Thinking Skills and Creativity*, 49, 101380. <https://doi.org/10.1016/j.tsc.2023.101380>
- Wang, R., Reis-Jorge, J., Crosta, L., Edwards, A., & Mudaliar, M. (2018, October 1–4). The use of social media and artificial intelligence tools by online doctoral students at the thesis stage. *Proceedings of Teaching and Education Conferences*, International Institute of Social and Economic Sciences, Vienna, Austria. <https://doi.org/10.20472/TEC.2018.006.006>

Copyright of Innovations in Education & Teaching International is the property of Taylor & Francis Ltd and its content may not be copied or emailed to multiple sites without the copyright holder's express written permission. Additionally, content may not be used with any artificial intelligence tools or machine learning technologies. However, users may print, download, or email articles for individual use.