Learned Publishing

ORIGINAL ARTICLE OPEN ACCESS

AI And the Editors' Ghost: Who Is the Writer Now?

David Clark[1] | David Nicholas[1] | Abdullah Abrizah[2] | John Akeroyd[1] | Jorge Revez[3] | Blanca Rodríguez-Bravo[4] | Marzena Swigon[5] | Tatyana Polezhaeva[6] | Anne Gere[7] | Eti Herman[1]

[1]CIBER Research, Newbury, UK | [2]Department of Library & Information Science, University of Malaya, Kuala Lumpur, Malaysia | [3]Department of Classics, University of Lisbon, Lisbon, Portugal | [4]Area de Biblioteconomía y Documentacion, Universidad de Leon, Leon, Spain | [5]Wydział Humanistyczny, Uniwersytet Warminsko-Mazurski, Olsztyn, Poland | [6]Library, Tomsk State University, Tomsk, Russia | [7]University of Michigan, Ann Arbor, Michigan, USA

Correspondence: David Nicholas (dave.nicholas@ciber-research.com)

Received: 12 September 2025 | **Revised:** 17 February 2026 | **Accepted:** 5 March 2026

ABSTRACT

This an exploration of the use of AI in research and writing. It builds upon the 'Harbingers' project, an international and longitudinal study of early career researchers (ECRs) and scholarly communication. In the fourth phase of the project, we returned to the theme of AI, in particular AI as 'ghostwriter'. Our sources are transcripts of conversational, open-form interviews with over 60 ECRs from Britain, Malaysia, Poland, Portugal, Spain, Russia, and other countries. For an initial analysis of the transcripts, we used Google NotebookLM. An overarching and thematic summary of the data was produced in minutes, that would otherwise have occupied our research team for weeks. The unprompted text, immediately plausible and coherent, was regarded by all national interviewers as impressive. Here, using a relatively small, convenience sample, we compare the AI generated summaries both against our original data and those first impressions. We reflect upon our own experience of using AI and that of our interviewees. This paper is about how we used AI as an experiment, our reaction to it, how that chimes, resonates, echoes the experiences of the ECRs. It is a calibration for our future data analysis.

1 | Foreword

The international Harbingers study of ECRs—their work life and scholarly communications—began a decade ago studying generational (Millennial) impact (H-1: 2015–2018); followed by pandemic impact (H-2: 2020–2022); and then in 2023, with the arrival of Large Language Models, AI (H-3: 2023–2024). Two years on, AI retains our attention (H-4: 2025). Following our observation of a 'ghostwriter in the machine', now we see growing application of 'AI' not just to writing, but also to search and even discovery.

Central to all Harbingers studies conducted over the last decade are the thoughts and actions of junior researchers—the professors of tomorrow—and their scholarly communications. Questions about career aims; assessment; reputation; metrics; general communication practices; information discovery; use; evaluation; trust; authorship; writing, publishing, and transformations are a common thread in all these studies of change.

2 | Introduction

The potential of AI to change scholarly communications was noted when Large Language Models (LLMs), notably 'ChatGPT', came to prominence in 2023. Our first study of AI in 2024 (Clark et al. 2025) identified aids to writing and translation as the most popular application ('the ghostwriter in the machine'). LLMs have developed considerably in the 18 months since that first study; the plausibility and coherence of the texts generated can be outstanding. Now, we are seeing applications to search and discovery, a capacity to abstract and generate a précis, that can seem both beguiling and unsettling.

Key Points

- Interview transcripts totalling 150,000 words rendered as a 1500-word summary in less than a minute.
- The summary rated overall 3/5 (passable), sometimes better by those who conducted the original interviews.
- A tendency to over-generalisation is the most observed fault.
- NotebookLM is a very useful tool to gain an overview, but best used as a guide for further scrutiny not as a draft.
- As we further explore the theme of 'the ghostwriter in the machine' we begin to discern there is also an editors' ghost.

A convenience sample of 62 ECRs from all subjects and 10 counties was assembled by local interviewers. They were interviewed for more than an hour and asked around 60 questions in a largely conversational format with many open-ended questions.

Wishing to form a first-cut overview in a timely manner, we used, as a speculative experiment, Google NotebookLM ("Your Personalized AI Research Assistant") to summarise over 150 thousand words of interview transcripts. Thus, we bring to this account an added dimension, our own somewhat unsettling experience of using generative AI. We set out to ask ECRs what they thought about using AI as a 'ghostwriter', and in the process observed it first-hand.

3 | Aims

Our aim was to look at the effect AI was having upon ECRs. To follow up on the 'ghostwriter' phenomena we had identified in earlier studies (Clark et al. 2025). It appeared then, that a major aspect of Generative AI would be the automation of 'wordsmithing' and a prospective 'Ghost Writer in the Machine'.

We present here a companion piece to a forthcoming analysis of the impact of AI on ECRs jobs and careers: an account of ECRs' use of AI, filtered through the lens of our own use of AI assistance to summarise and give account of our findings. An overview that in previous Harbingers studies we would venture only after many months' exegesis of the transcripts and quantitative summation.

There has been an extraordinary expansion of LLM-based AI since 2023 and while we remain cautious as to the results, it is apparent that the potential of the tool cannot be ignored. We append an example of NotebookLM output unedited: it is, in the judgment of those closest to the original data—our interviewers—at least a passable account, occasionally perceptive; but a final judgment must await those many months of study.

4 | Working Definitions

4.1 | Early Career Researchers

There are various, conflicting, and country-specific definitions of an ECR (Teixeira da-Silva 2021). Our definition is pragmatic: employed in a research position, relatively young, in an early career phase, not yet established as permanent faculty.

4.2 | Artificial Intelligence

Artificial Intelligence has no firm or formal definition. Consequently, when asking about attitudes to and anticipations of the place of AI it is necessary to consider—'what do you mean by AI?' We have previously found that even quite unremarkable word processing tools were considered 'AI' by some of our respondents (Clark et al. 2025). The scope of 'AI' here is inevitably biased by a contemporary interest in Generative AI and Large Language Models, thus only a sub-set of what may be considered as Artificial Intelligence. We are thus talking about the capacity to generate a plausible text—a convincing and realistic fiction. The process is stochastic; the output text is plausible because the words are arranged in ways that match the probabilities derived from the training input. But the process is occult; there is limited facility to determine, verify, or analyse the logic or 'reasoning' of the process. Contrast that with symbolic AI (McCarthy et al. 1955; McCarthy 1958), autonomic and mathematical models where there is an 'audit trail'—an algorithm that may be verified or an activity that can be tested and demonstrated to work. Perhaps the advent of 'reasoning' models may change that; we think it too early to say. With Generative AI we can only judge by what is presented: truth or fiction?

5 | Literature Review

The development of LLMs and AI in the past 2 years has been rapid; the sometimes-ephemeral social response to this technology is often too fluid to be memorialised in 'peer reviewed literature'. We found only one recent study that had directly tested AI-generated academic summaries for factual accuracy (Erol et al. 2025); others (Corbin and Walton 2025) have proposed a 'research agenda'. Farber (2024) considers the uses of AI in the reviewer selection.

Two earlier *Harbinger* studies, published in this journal, contain extensive literature reviews covering the period before 2025 and should be consulted (Herman et al. 2024; Nicholas et al. 2024). In general, we found that ECR focussed studies found evidence that junior academics had more interest and positive views of the technology than senior faculty, and more of them had already tried it. A relevant study published since (Kim et al. 2025) carries a similar message, finding that late early-career to early mid-career researchers tend to adopt AI more actively than others, although this pattern varies across disciplines. Lastly, Perplexity.ai, an AI enhanced search engine, provided similar conclusions when asked ‘Are ECRs taking to AI for their scholarly communications’. After examining 20 sources (mostly ones we had published or which we had already identified) it came back with the overarching ‘key findings’ that there was strong adoption, active experimentation and high usage rates by ECRs.

6 | Methods

6.1 | Recruitment of Interviewees

National interviewers (from Malaysia, Poland, Portugal, Russia, Spain, UK and US) recruited ECRs, using their local research networks and from previous iterations of *Harbingers* (32 were recruited this way). Each country was allocated a quota of interviewees (10–11). Malaysia, Poland, Portugal and Spain did recruit 10 interviewees, and Russia 11, but resource constraints led to a merged UK/US ‘international’ group including a British ECR working in Ireland, a Romanian working in the Netherlands, an Australian who had worked in the UK, a Canadian working in the UK, plus a German and an Italian ECR. This international group consisted of 11 ECRs.

With limited resources, this was very much a convenience sample; our interviewers' capacity to balance age, gender, and subject was limited. The overall composition had several irregularities, so any demographic partition of the results needs to be treated with caution. We have interviewed *people*, not categories, and the phenomena are developing fast; any generalisation is provisional, underpinned by a recognition that what we observed is consistent with our own experience.

A breakdown by country, discipline, gender, and age-band is given in Table 1. Note especially the age distribution, a median age of 34, with one in five ECRs being 40 or over; it could be argued that these are unlikely ECRs. This is part-explained by ECRs retained from previous stages of the project, all a year or two older. But it is also a reflection of the *precarity* of academic employment: not everyone moves forever upward and tenured jobs are difficult to obtain. We had more women than men and more scientists. Scientists proved more interested in a topic relevant to many of them.

6.2 | Data Collection

Semi-structured, free-flowing, conversational interviews of 60–90 min in duration were the main source of data. The interview schedule contained questions on general scholarly communication attitudes and practices as well as specific questions about the role of AI within scholarly communications. These general questions provide us with continuity with our previous Harbingers studies: a background and context for the focus on AI in this round. We list here only the questions about AI (Table 2).

6.3 | Data Analysis

All interview transcripts, having been read and approved by the interviewees, and translated to English where necessary, were transferred by the national interviewers to a coding sheet. This structured the transcript in a tabular form, with codification of responses to suitable questions as Yes/No/Uncertain.

TABLE 1 | Demographic breakdown of ECR convenience sample in 2025.

	Discipline									
	CHEM	**ENVIR**	**HUM/ARTS**	**LIFE**	**MATH**	**MED**	**PHY**	**SOCH***	**SOCS****	**Total**
N	4	5	12	4	5	5	5	7	15	62
%	6	8	19	6	8	8	8	11	24	100

	Country						
	ES	**GB/US/Int*****	**MY**	**PL**	**PT**	**RU**	**Total**
N	10	11	10	10	10	11	62
%							100

	Age					
	Youngest (26–30)	**Younger than most (30–34)**	**Median (35–37)**	**Older than most (37–39)**	**Oldest (40–51)**	**Total**
N	12	13	12	13	12	62
%	19	21	19	21	19	100

	Male	**Female**	**Total**
N	27	35	62
%	44	56	100

Note: * Includes Economics and Business, Geography and Psychology. ** Includes Anthropology, Philosophy, Politics and Sociology. *** This is really a composite international grouping including a British ECR working in Ireland, a Romanian working in the Netherlands, an Australian who had worked in the UK, a Canadian working in the UK German and an Italian ECR.

TABLE 2 | interview questions on AI and scholarly communications.

Broad topic	Question asked
Reputation	Can AI help in reputation building?
Skill development	Have you tested or considered AI-based tools for research work? What may be the benefits or problems involved in using AI for research?
	Are AI-based tools a regular part of your research work?
	Any training or guidance in application of AI by their institution?
	How far does reliance on technology weaken fundamental skills? Can developing new technological skills compensate for this loss?
Information discovery, use, and trust	As shown by Google Gemini etc. search results are moving to presenting an AI-generated summary of results. We know that general users rarely scroll-down the results, do academics take more care? What steps, if any, do you take to follow-up, examine, cross-check, and generally verify what you find online?
	Where do you go to search for formal scholarly communications, such as peer reviewed articles? If that fails where next? Has the introduction of AI enhanced search and presentation changed the way you do this?
	What of the capability of AI to summarise and translate research? Do you use AI-based tools for this purpose? If so, how do you follow-up, examine, cross-check, and generally verify an AI-generated result?
	When you have searched and found a published article on a topic important to their research, what criterion persuades you to read it?
	To what extent do you think that the peer review system vouches for the quality and trustworthiness of formally published research? Has the introduction of AI to the peer review system changed this?
	Do you think that the peer review system needs improving in any way?
	How do you decide whether to trust informally disseminated results in your own specialisms? Has introduction of AI changed your criteria?
	What would make you suspect that published material was possibly AI generated?
Authorship, publishing	Have you used AI to make a first draft, recast or rewrite your own (original) text.
	AI tools, just like colleagues, can suggest topics for new research, by providing insights, criticism and feedback—Have you used AI tools as a prompt or sounding board (validator) to develop or explore an idea?
	Have you used AI-based tools to assemble data sets and analyse them?
	Generative AI can expedite the manuscript development process, particularly for non-native speakers of English. May produce an entire research paper. This raises these questions: Ownership—to whom should the text belong? Control—who has control over content? Leadership—who takes the lead in writing?
	"AI will be like newspapers; when you read about something you know, you can see what is obviously wrong, but when you don't know you believe it anyway" Do you agree?
	Are authorship policies likely to change because AI has the potential to become a 'ghostwriter' (another author)?
	Machine learning is only as good as the data used to train it. Thus, large language models may have embed biases, assumptions etc. Do you think these are essentially different in scope or kind to similar limitations in conventionally published research material?
	AI could increase productivity but lower quality; how is scholarly integrity to be maintained in these circumstances?

(Continues)

TABLE 2 | (Continued)

Broad topic	Question asked
	Are the interaction of 'Open Access' and 'predatory' publishing practices 'poisoning/polluting the well' of source material used to train AI?
Outputs—integrity, Ethics	Do you believe that the AI-associated potential for rapid production of scientific articles leading to a decline in the overall quality of research output?
AI general	Awareness of AI: Experience and Encounters; Engagement; Utility; Reservations.
	Are you concerned about the use of AI in any way? If so, why?
	If current AI developments bring big changes, how soon: now, in a year or two, this decade, in your working lifetime? Ethical implications?
	Do you believe that AI is introducing new challenges related to scholarly integrity and ethics?
	If AI is seen as a transformative force, in what ways will it reshape academic practices, research, and knowledge dissemination?
	How do you navigate the tension between the concept of 'publication' (sharing knowledge openly) and 'copyright' (protecting intellectual property) when it comes to using large datasets for training AI models?
	Anthropomorphism: the tendency to give conversational interfaces to AI (e.g., chat-bots) encourages the perception of human-like personalities Is this making things 'user-friendly' or betraying a naive trust

Comments were also added with information derived from additional enquiries or clarifications. Thus, the coding sheets held quantitative and qualitative responses to questions, but primarily direct quotations with occasional explanatory comments from the interviewers.

The 'codesheet' was developed for Harbingers-2 (circa 2021) as a means of enabling an overall view and quantitative summary of transcripts amounting to ~800,000 words in total. The processing was refined during that project, the information gathered from documents and spreadsheets being restructured as an SQL database which allowed for full-text searching and further analytic partitioning of the data. This enabled a shift in our methodology, away from statistical aggregates based on the coding of answers to a deeper consideration of the transcripts and of individual ECR voices. But this was very much a hand-crafted search of the texts; could we employ some more effective, automated techniques? In the spring of 2023, our own internal assessment of the potential was cautious:

> For our small sample sizes it would likely always be more reliable to just skim though the text by someone familiar with the project and our context and interpretation of the topic. The overhead of set-up and training would be a far greater effort with a result that would still be subject to doubt.

A lot has happened in 2 years, and anecdotal opinion suggested the technology might now be effective and affordable for our purposes. Speculatively, we uploaded the quotes and interviewer comments to Google's NotebookLM and in seconds, unprompted, it produced a summary. This use of AI for preliminary analysis of the interview data, appeared useful in identifying main themes, frequent terms and consensus. But we were uncertain of its depth and capacity for discernment. Does it help or hinder the interpretation of cultural and discipline-specific context, not to mention the muting of rhetorical and colloquial figures-of-speech? Perhaps our interviewers would be better able to detect this and other flaws in a summary consisting of only *their* transcripts; those interviews of which they possessed first-hand and unique knowledge.

So, we tried again, this time with a separate upload of each interviewers' data. This time with the full text of quotes, comments, and where relevant, encodings, plus a statistical summary of the coded responses. For each set of interviews data we set NotebookLM two tasks. First a 'Briefing Document', a one-click option generating a summary of the sources with copious references thereto. Secondly, a prompt to focus on a topic of particular interest: "Do the ECRs interviewed here consider Artificial Intelligence (AI) as, on balance of probabilities, an opportunity or a threat to their future career?" (This report is the prime focus of our companion paper: Do ECRs consider AI as an opportunity or a threat?).

For each report we asked our interviewers to look specifically for anomalies, inaccuracies, and that subjective 'uncanny valley'—the intuitive sense of something not quite right (Mori et al. 2012). We also sought an overall impression of how good, how plausible, each report was. As a guide we suggested a broad rating schema (see below) with a proviso that this should be assessed relative to an intended audience, in our case: was this at least good enough for post-grad work?

*	= fail	obvious errors (contradiction, not 'common-sense' etc.) Even a non-expert can see flaws.
**	= poor	errors noticeable by someone familiar with the subject.
***	= passable	'coherent & plausible' No apparent errors or nonsense but bland, generic, non-specific, a consensus. Nothing you do not already know. Possibly beguiling and superficially persuasive.
****	= perceptive	Some original observation, dug a little deeper could seem five-star to the non-expert.
*****	= understanding	analytic, beyond the surface, real insight.

7 | Results

Our use of NotebookLM was, and remains, speculative: it was not an intentional component of our methodology for Harbingers-4. The NotebookLM texts proved to be quite stunning in their plausibility and coherence, but leave a suspicion that it is all a bit too smooth and bland (Appendix A). The NotebookLM reports tend to an equivocal 'on the one hand', 'on the other hand' style. But they certainly provide food-for-thought, an outline, a pointer to what needs a closer look. As such they may supplant our aggregation of encodings as a convenient and apparently easy way to assess an overall weight-of-opinion. They also give us another insight; we have experienced at first-hand what our ECRs tell us. The Harbingers studies have always been a case of academics studying academics; in this case, we have become users of AI studying users of AI.

How should we evaluate these AI generated summaries? The original source material is subject to several layers of interpretation, of editorial intervention. Orality is not literacy; the transcribed spoken word is tidied-up for interviewee approval. They may add corrections and clarifications. In most cases, the approved transcript will have been translated to English. Might we devise some independent criteria, an assessment by an independent observer? Our resources are limited; the interviewer is also transcriber, translator, interpreter: the only first-hand authority.

What follows are the comments of *our interviewers* (sometimes citing individual ECRs) after looking at reports on *their* interviews. We divide them into sub-sections headed by key themes. The overall rating for every one of our interviewers was at least 'passable' (***) with most rating it a little higher although with some hesitation in awarding an unequivocal 'perceptive' (****). The rating of the briefing report with its extensive citation back to the provided sources tended to be a little higher than that of the prompted question. This may be due to our lack of skill in 'prompt engineering'—knowing how to frame the right question—and what appears to be some inconsistency in the way citations are presented, the displayed on-line pop-up citations not being present in the offline copy of the text. This is a new tool for us; we have much to learn. This is the reaction of our interviewers:

> I've gone through the findings, and overall, they look great, the analysis checks out well against the codesheets and reflects the data accurately.

> The quality of the Google AI tool is astonishing. It seems to me that qualitative analysis tools like ATLAS.ti and others are forecasting their final days.

> I have gone through the text and I think I am leaning towards *** at least—in that it is hard to see that any of this is in anyway "incorrect"

> That said, the summarising ability is impressive. [...]. My evaluation clearly falls between *** passable and **** perceptive.

7.1 | Unwarranted Generalisations

The most common observation was of overgeneralisation. The sources were restricted to that of a single interviewer, so that first-hand knowledge could better detect any misrepresentation. This overgeneralisation may therefore be but an artefact of our methodology for this test. It that case it might be expected to disperse with the full set of source material.

> ...it seems to me that there are sometimes unwarranted generalisations. I'll give an example: I only found one reference to job replacement, yet the output states:

- *Job Replacement and Dehumanisation*: *Some express a direct fear of being replaced by AI.*

> This seems exaggerated.

> Another example that also strikes me as an overgeneralisation—since I only found a single reference to it—is:

- *Peer Review Challenges: While some see AI as potentially improving peer review efficiency.*

> A third example is the mention of "*Broader concerns include AI's potential for societal manipulation (e.g., political campaigns)*," when in reality the reference is fairly superficial ("*Just look at TikTok and the various political and anti-government campaigns orchestrated by AI*.")

> That said, the summarising ability is impressive. Perhaps due to the small sample size, these generalisations occur more easily.

> It's about generalisation, overgeneralisation.

> For example:

> J*ob Displacement and Loss of Creativity*: One ECR specifically mentions knowing people who lost their jobs (e.g., illustrators, copywriters) to AI, fearing it will take away creative jobs and replace them with "slop." Others fear it may "*doom creativity*" of future generations.

7.2 | 'Briefing Doc' Superior to Prompted Response

> Regarding the briefing document, the writing style is very similar to what we have used in previous project publications, with the inclusion of quotes from the ECRs. I believe it can clearly be rated as **** = perceptive, as it provides a well-supported reading of the interview findings.

> The briefing document is correct—**** perceptive. There are references to quotes from statements everywhere, it looks very good!

7.3 | Style, Clarity

> There are issues to do with style (grating to the point of being unreadable) and—in my view—an overuse of quotes or maybe of quotation marks. The bullets seem to have a sameness about them which is probably as it will be and why some were saying they can easily tell AI generated text.

7.4 | Weighting and Emphasis

> It is hard to see that any of this is in anyway "incorrect"—if there is an issue it is more about what weight the machine gives to a view expressed which in turn is a judgement on the perceptiveness/validity of the interviewee which only the interviewer can make. I assume the AI engine accepts all input as equivalent. You might somehow want to weight each interview and in turn convey that to the engine. Though I guess across the whole cohort the AI will take some account of frequency.

7.5 | Participant Voices

Then there is the status of 'participant voices' Sometimes the AI summary rephrases the text so it becomes estranged: For example, this generated text "Job Market Disruption: There is a concern about the imminent disruption of the job market, with AI potentially making certain professions or junior roles obsolete." appeared to the original interviewer to be an 'AI hallucination'. Yet, with scrutiny we can uncover a basis for this in the transcripts: "*It's advancing at breakneck speed. Within just 2 years, we might already feel its disruption in the job market—certain professions could become obsolete as AI takes over their functions. Frankly, I believe this shift is imminent*." Is this a reverse 'hallucination', something overlooked by the interviewer but nonetheless present in the transcript? Or, perhaps lost in translation: from the original interview, translated to English, frozen as text rather than remembered conversation, mangled again by the LLM? AI has captured a form of words but lost something of the context and possible sub-text.

Another observation on job-losses, this time by an ECR, introduces another perspective on participant voices:

> …they've already said that psychologists' work is now shifting to artificial intelligence because, in many ways, AI is better at holding conversations—completely so. I was just playing around here, sitting and asking it questions, like what to do and how. And it really does come up with such well-structured phrases, sentences, and options—it just does everything perfectly. And it's much faster, in my opinion. But many people are very closed off; they're afraid to express their opinions to another person, whereas here, you can say everything openly.

In this case we see an echo of Weizenbaum's observation of the responses to his ELIZA programme (today we would call it a Chatbot) more than half a century ago:

> I was startled to see how quickly and how very deeply people conversing with [ELIZA] became emotionally involved with the computer and how unequivocally they anthropomorphised it. Once my secretary, who had watched me work on the program for many months, and therefore surely knew it to be merely a computer program, started conversing with it, After only a few interchanges with it, she asked me to leave the room.

> [...] What I had not realised is that extremely short exposures to a relatively simple computer program

> could induce powerful delusional thinking in quite normal people.
>
> (Weizenbaum 1976).

7.6 | The Editors' Ghost

We set out to explore the observation we made in Harbingers-3, the use of AI as a ghostwriter (Clark et al. 2025). What is emerging here, a preliminary result of our own experiment in using this new tool, is the power of using LLM-AI as a research assistant. An easier way to sift through piles of data, to classify, aggregate, induce some hint of order in the chaos. The risk here is that without real, live intelligence, we cannot see the wood for the trees, the editorial function of knowing what counts, and what is worth counting, is lost. In sum, the editor becomes a mere shadow, a ghost. As one of our team observed:

> AI tools seem to be good at summarizing, structuring text, selecting quotes. But they miss interesting moments in interviews, such as the respondent's attitude toward technology in general. For example, I was interested to learn that one of the researchers combined logical analysis in his work with an almost shamanistic attitude toward technology.

The value of the extended interview format is that sometimes we can pick out one observation of real significance. As we have developed the Harbingers project, we have moved away from aggregates—so many percent said this, so many said that—the individual voices count, not how many are counted. And thus, what seems the fundamental weakness of generative AI: the absence of perception, of seeing something worth noticing, not humming along to the common tune.

This is not just our impression; it is something our ECRs also see.

> AI really loves writing bullet-pointed or numbered lists—it's an immediate formal marker. There's also a lack of what I'd call 'subjectivity;' no substantive critique of the content. If the AI does raise content-related questions, they're phrased as suggestions like, 'I'd recommend changing this to that,' rather than outright declaring, 'This is all wrong.'

> ...finding literature, analysing large volumes of text, and extracting only what's relevant to a specific study—it's incredibly convenient and, of course, a huge time-saver.

> But on the other hand, as a researcher, I can say that reading literature—not just on my exact topic, but also works where I hoped to find something relevant—has repeatedly led me to discover ideas that sparked new directions in my research. Sometimes, it was information completely unrelated to my topic, yet it prompted me to think differently about my subject.

> So while the purely pragmatic approach AI enables is undeniably useful—yes, it's efficient—it also, in my opinion, risks dulling the researcher's mind. That mind should remain active, engaged in the investigative process. After all, research isn't just about structured analysis—it's also about those moments when you're walking down the street, observing the world, and still turning over ideas related to your work.

> ...it's undeniably useful for discovering relevant papers. Run a script, and it can surface high-impact journals or instantly pull up the 20 most relevant articles for your keywords—no manual searching required. That's a clear plus. But if you rely on AI to summarize instead of doing it yourself? That's a downside. You risk missing nuances or failing to deeply engage with the material.

7.7 | Plagiarism

Another theme often touched upon in both NotebookLM's summaries and the interviews was that of plagiarism. There is a risk of unintentionally copying phrasing or ideas from AI's training data if ECRs do not properly edit or acknowledge AI use. But is it a case of copying 'from AI's training data'? There is a nuance here that needs explication. Some might argue that 'the AI model' is the plagiarist in using original material as training data. Or, it could be argued that 'the AI model' has been created by absorbing and processing the training data; it can no more be accused of strict plagiarism than can any diligent student who has read and absorbed the literature. Or, is it plagiarism of 'the AI model' by those who quote it without due acknowledgement? As one of ECR put it:

> The text certainly belongs to the author, who must maintain control over its content. Should AI be listed as a co-author? No, it shouldn't. After all, we don't credit the pen as an author, nor the keyboard model or motherboard inside the computer. However, in the methodology section of a paper, one could mention that the text was drafted with AI assistance. I have an inarticulate sense of why this matters, though articulating it proves difficult—largely because AI writes so convincingly that it feels like there's a human behind the screen. If it can write, doesn't that make it human-like? Yet rationally, we know it's just an algorithm. Much of its output barely differs from human writing, especially for the routine task of transposing thoughts into on-screen text.

> [...] my stance is conflicted. On one hand, these are my ideas; on the other, I welcome their reuse—even via AI—to advance my field. Why not? Proper attribution suffices: a mandatory historiographical review crediting prior work. But claiming sole ownership is rarely straightforward. [...]. My 'originality' lies only in combining these elements. Can I truly call the resulting synthesis 'my' idea? Unlikely. My thoughts are already dissolved in the collective stream of literature, conversations, and readings that shaped them.|

We can think of this another way: if we feed our interview data into NotebookLM and it produces a report—a report that appears to be a reasonable overview of the material—are we right to claim that this is our data, that automatic processing is no different in kind but only by degree, equivalent to using software to compile a statistical analysis? If I use a statistical function with little knowledge of the underlying principles, can I claim the result as all my own work?

7.8 | Prompt Engineering

The advent of LLMs has certainly advanced the concept of a natural language interface to computing, but we doubt that this means anyone can make use of it. It seems more likely that we will see the emergence of 'prompt engineering' as a new specialist skill. The productivity of the best programmers and other users is enhanced by the potential of automation, but offset by the growth of 'sorcerers' 'apprentices', given powers beyond their skill and comprehension. As an ECR observed:

> Everyone complained when articles started going digital, saying it meant researchers didn't know how to properly search for literature. But here's the thing: every new technology demands new skills, and eventually, everyone adapts. There are always pros and cons—the key is identifying the advantages and leveraging them.

A history of programming languages could be written round that theme: make them more like natural language, the supposition being that this would make programming something anyone could do. But the result tends to be a new specialist language, it enhances the performance of those who master it. As our ECRs find:

> Through trial and error, you gradually learn how to craft prompts that deliver the results you want. Last winter, I noticed more detailed popular articles emerging about proper prompt engineering (as they call it). I've picked up techniques from these—like instructing the AI to adopt specific personas or professional roles to improve responses. My knowledge comes mostly from hands-on experience and these accessible guides rather than comprehensive manuals.

8 | Conclusion

It may be that we, novices at 'prompt engineering', are not yet very good at framing the right questions (prompts), but our experience suggest that NotebookLM's preset one-click report, the 'Briefing Document' is a surer option in that it produces a document that gives detailed citations of the sources provided and is thus much more reliable as a basis for refinement. By contrast it appears that some caution and developed skill at the prompt is required if the objective is to develop in-depth understanding as opposed to a fluent, but potentially flawed analysis. Another observation is that providing both transcript and interview schedule may introduce a bias. For the LLM all source text is 'grist to the mill', we suspect that some of the observations for which we cannot find any evident response arise from NotebookLM not distinguishing what are in effect leading questions from responses.

Overgeneralisation or lack of specificity? The problem is not so much that it does not pick out the unique and particular, but that of recognising the qualities that make something special and notable. The generated text looks very plausible, but cannot be trusted as a reliable summary of sources. Overgeneralisation from limited material in the source was noted by several of our interviewers. That could become even harder to spot once many interviews and sources are combined.

In earlier stages of Harbingers, we made much use of coding and statistical aggregation to render a mass of interview transcripts into a coherent summary. But that means either posing only apparently simple (yes or no?) questions or losing depth and insight by fitting to a Procrustean bed of presumed categories. We have moved away from that method toward attending to individual voices, using our judgment to discover an account worth attention. In the past 2 years LLMs have advanced sufficiently that they can be used to aggregate individual voices direct from the transcript: it seems we can omit the coding and go straight to the overview—a nice smooth text that sums it all up. But the trade-off is that we risk a loss of detail, the telling individual insight. The result is an aggregate of what our interviews record, filtered through an aggregation of the language models' training data: processed peas turned to mush?

To ordain that AI cannot be an author is to say AI must be a ghostwriter, that is, one never acknowledged. In transcribing, translating, summarising, we necessarily introduce our own perceptions and interpretation, and thus, if we use AI, a ghost editor. This journal, being for practitioners, is not the venue for a philosophical discussion, yet we cannot conclude here without noting a deeper dimension: 'the ghost in the machine' (Ryle 1949), that unease that comes when we discover yet another task, supposedly uniquely human, that machines can do.

For illustration in this article, we have generated a 'Briefing Doc' by Google NotebookLM (see Appendix A). It differs from our earlier versions in that, for reasons of privacy and confidentiality, we have omitted demographic data that may occur in the quotations and comments.

Author Contributions

David Clark provided oversight of the article, analysed and wrote much of it. David Nicholas edited and commented on it. The rest of the authors undertook their country analyses, insights, and provided general feedback on the paper.

Acknowledgements

The project was largely self-funded by members of the Harbingers group: CIBER Research, University of Leon, University of Malaya, University of Warmia and Mazury in Olsztyn, Tomsk State University, and University of Lisbon. Thanks also to Anne Gere at the University of Michigan.

Consent

Participants who freely opted to take part in the interviews were asked to provide their names and contact details for follow-up questions regarding the accuracy of the interview transcripts, but access to all personal data was restricted to the investigating team and was removed before the analysis of the results.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

Clark, D., D. Nicholas, M. Swigon, et al. 2025. "Authors, Wordsmiths and Ghostwriters: Early Career Researchers' Responses to Artificial Intelligence." *Learned Publishing* 38: e1652. https://doi.org/10.1002/leap.1652.

Corbin, T. A., and J. Walton. 2025. "The Missing Story of GenAI Summarisers: A Critical Research Agenda." *Higher Education Research and Development* 44: 1–14. https://doi.org/10.1080/07294360.2025.2486185.

Erol, G., A. Ergen, B. Gülşen Erol, et al. 2025. "Can We Trust Academic AI Detective? Accuracy and Limitations of AI-Output Detectors." *Acta Neurochirurgica* 167: 214. https://doi.org/10.1007/s00701-025-06622-4.

Farber, S. 2024. "Enhancing Peer Review Efficiency: A Mixed-Methods Analysis of Artificial Intelligence-Assisted Reviewer Selection Across Academic Disciplines." *Learned Publishing* 37: e1638. https://doi.org/10.1002/leap.1638.

Herman, E., D. Nicholas, A. Abrizah, et al. 2024. "The Impact of AI on the Post-Pandemic Generation of Early Career Researchers: What We Know or Can Predict From the Published Literature." *Learned Publishing* 37: e1623. https://doi.org/10.1002/leap.1623.

Kim, S. Y., W. K. Lee, S. J. Jee, and S. Y. Sohn. 2025. "Discovering AI Adoption Patterns From Big Academic Graph Data." *Scientometrics* 130: 809–831. https://doi.org/10.1007/s11192-024-05228-4.

McCarthy, J. 1958. "Programs With Common Sense in NPL." In Symposium on the Mechanisation of Thought Processes. http://jmc.stanford.edu/articles/mcc59/mcc59.pdf.

McCarthy, J., M. Minsky, N. Rochester, and C. Shannon. 1955. "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence." http://jmc.stanford.edu/articles/dartmouth/dartmouth.pdf.

Mori, M., K. F. MacDorman, and N. Kageki. 2012. "The Uncanny Valley." *IEEE Robotics and Automation Magazine* 19, no. 2: 98–100. https://doi.org/10.1109/MRA.2012.2192811.

Nicholas, D., M. Swigon, D. Clark, et al. 2024. "The Impact of Generative AI on the Scholarly Communications of Early Career Researchers: An International, Multi-Disciplinary Study." *Learned Publishing* 37: e1628. https://doi.org/10.1002/leap.1628.

Ryle, G. 1949. *The Concept of Mind*. University of Chicago Press.

Teixeira da-Silva, J. A. 2021. "Challenges That Early Career Researchers Face in Academic Research and Publishing: Pre-and Post-COVID-19 Perspectives." *Exchanges: The Interdisciplinary Research Journal* 9, no. 1: 77–106. https://doi.org/10.31273/eirj.v9i1.882.

Weizenbaum, J. 1976. *Computer Power and Human Reason: From Judgment to Calculation*. W.H. Freeman.

Appendix A

The 'Briefing Document' below was generated by Google's NotebookLM. It differs from earlier versions quoted in the main text in that we have removed some personal data that may be present in our original source text. It should be read as a fair example of the output generated by NotebookLM; its accuracy and completeness are untested, nor can we vouchsafe it is free from 'hallucinations'. All we can say, with reasonable confidence, is that like our other experiments it is on the face of it a plausible account of the data we have gathered.

It is derived from four 'sources' uploaded to the LLM: the full-text quotations, coding and comments from the interviews (over 150,000 words); a list of the questions, a brief description of each ECR; and a compilation of tables analysing the encoded questions (e.g., demographics).

The document is presented 'as is' with one exception. Extensive references to our sources in the form of keys to ECR and Question have been removed and replaced with a count of the total number of such references. For example, {cited 0}. Absent the source documents these references serve no purpose here. This has removed much clutter from the text and preserved confidentiality of our sources. Note that the use of bold text to highlight text in the prompted follow-up section is entirely the work of NotebookLM.

Briefing Document: The Impact of AI on Academia and Research

This briefing document synthesises key themes and facts from the provided sources regarding the current and future impact of Artificial Intelligence (AI) on academic practices and research.

Perceptions and Adoption of AI in Academia

There is a widespread, though varied, adoption and perception of AI tools among academics. While some are enthusiastic "techno-optimists" {cited 1} who embrace AI's potential, others maintain a cautious or even skeptical stance.

- Growing Use and Awareness: Many researchers are aware of and increasingly using AI tools. ChatGPT, Gemini, and DeepSeek are frequently mentioned for tasks like drafting content, summarising papers, and brainstorming {cited 6} Some even report using AI almost daily {cited 2}.
- Curiosity and Experimentation: Many are driven by curiosity to explore AI's capabilities {cited 3}. Training courses, where available, are seen as useful for "open[ing] [their] mind" {cited 1} to potential applications.
- Scepticism and Reservations: A significant number of academics express scepticism, particularly regarding AI's ability to perform complex, nuanced tasks or generate truly original content {cited 3}. Some believe that "AI is already being overblown. It's a complex analysis tool, right? And then it can spit it back out at you. So, anything that happens is it just makes your skills more efficient" {cited 1}.
- Informal Learning: Many researchers, particularly in fields where formal training is lacking, learn to use AI through "trial and error, learn online, and tips from colleagues" {cited 3}.

Benefits and Utility of AI

AI is widely recognised for its potential to improve efficiency and streamline various academic tasks, freeing up researchers' time for more complex work.

- Enhanced Productivity and Time-Saving: AI tools are lauded for "speeding things up" {cited 1}, particularly in tasks such as literature reviews, summarising articles, drafting proposals, and checking grammar {cited 6}.
- Writing Assistance: AI is extensively used for "improving the writing of research outputs" {cited 1}, rephrasing text, correcting grammar, and overcoming language barriers for non-native English speakers {cited 4}.
- Data Analysis and Research Exploration: AI can "process and comprehend large amounts of data into something that's a bit more digestible" {cited 1}, "uncover hidden patterns in data" {cited 1}, and suggest research directions or interdisciplinary connections {cited 3}.
- Idea Generation and Brainstorming: Many find AI useful for "brainstorming ideas" {cited 1}, "formulate research questions" {cited 1}, or even "test ideas" {cited 1}.

Concerns and Challenges

Despite the perceived benefits, significant concerns surround the ethical implications, quality control, and potential negative impacts of AI.

- Erosion of Critical Thinking and Skills: A major worry is that over-reliance on AI may "reduce each person's critical thinking" {cited 1} and "weaken fundamental skills" {cited 1} such as deep analytical skills, rigorous research, and even manual coding {cited 3}.
- Misinformation and "Hallucinations": AI's tendency to "invented" bibliographical citations {cited 2}, fabricate content {cited 1}, and generate "fake news" {cited 1} is a serious concern, leading to a "growing distrust, the need for constant verification" {cited 1}.
- Plagiarism and Authorship Ambiguity: The ease with which AI can generate text raises fears of widespread plagiarism among students and researchers {cited 2}. There is a lack of clear guidelines on "where do we draw the line between human and machine contributions?" {cited 1} and whether AI should be acknowledged as an author {cited 2}
- Bias Amplification: AI models, trained on existing data, can "amplify existing biases" {cited 1} from their human trainers or the datasets themselves {cited 2}. These biases can be harder to detect in AI-generated content {cited 2}.
- Quality vs. Quantity: There is a concern that AI may lead to a "decline in the quality of research" {cited 1} by facilitating "rapid production of scientific articles" {cited 1} without sufficient rigor or originality.
- Intellectual Property and Data Usage: Questions are raised about the "intellectual property rights" {cited 1} of data used to train AI models, particularly copyrighted and personal data {cited 2}. Many argue for transparency and consent in data usage {cited 3}.
- Security and Privacy Risks: Attaching sensitive documents to AI applications raises data protection concerns {cited 1}.
- "Humanising" AI: While some believe "humanising" AI fosters acceptance {cited 1}, others warn that it can lead to "naive trust" {cited 2}, with users potentially overlooking its limitations and the non-human nature of the tool.

Academic Practices and Evaluation

AI is influencing traditional academic processes, from research dissemination to peer review and career progression.

- Research Dissemination: Researchers use platforms like LinkedIn, ResearchGate, and academic social networks to disseminate their work {cited 4}. AI can assist in "design[ing] a promotional plan" {cited 1} or crafting engaging social media posts {cited 2}. However, some fear that AI's impact may lead to a "loss of credibility" {cited 1} for informal media.
- Peer Review: The peer review system is seen as "fundamental, albeit imperfect" {cited 1}. AI's potential role in speeding up checks for plagiarism, citation verification, and formal language review is acknowledged {cited 2}. However, many insist that "human judgment is still crucial" {cited 1} and that AI should not replace expert

evaluation. Concerns are also raised about reviewers using AI themselves to generate feedback {cited 2}.

- Assessing Research Quality: Citations remain a key metric for academic reputation and evaluation {cited 4}. However, there is growing recognition that they "don't always adequately reflect the value or impact of research" {cited 1}. Alternative metrics, such as social media mentions, are gaining importance for measuring social or communicative impact {cited 2}.
- Impact on Publications: The value and perceived quality of traditional publications might change {cited 2}. Some believe publications may shift to a "model that only features results, discussion, and little else, because the rest is done by AI" {cited 1}. Conversely, others hope AI might "enhance the value of those used responsibly and transparently" {cited 1} by making it easier to identify novel work.
- Predatory Journals: The proliferation of predatory journals and questionable publishing practices is a significant concern {cited 3}. AI's ability to generate low-quality content could exacerbate this issue, "poisoning the well" {cited 1} of training data for future AI models.

Future Outlook and Recommendations

Academics anticipate significant and rapid changes due to AI, with many expecting major transformations within the next 5–10 years {cited 3}.

- Need for Guidelines and Policies: There is a strong consensus on the need for clearer ethical guidelines and institutional policies regarding AI use in research and teaching {cited 4}. These should address authorship, data transparency, and responsible implementation.
- Emphasis on Human Oversight: The role of the human researcher as the "final say" {cited 1}, the "writer" {cited 1}, and the "leader" {cited 1} remains paramount. AI should be a "support tool" {cited 1} to "amplify our impact"{cited 1}, not a replacement for human intellect and accountability.
- Continuous Learning and Adaptation: Researchers and institutions need to adapt to these changes, develop new technological skills, and integrate AI in a balanced and thoughtful manner {cited 2}.
- Prioritising Quality and Integrity: Maintaining scientific integrity by valuing quality over quantity {cited 2} and rigorously verifying AI-generated content is crucial to navigating the evolving landscape.
- Open Access and Data Responsibility: The principles of open science and ethical data use are highlighted, advocating for open knowledge and proper attribution for data used in AI training {cited 2}.

In conclusion, AI is rapidly reshaping academia, offering powerful tools for efficiency but also presenting complex ethical and practical challenges. The academic community is grappling with these changes, seeking to harness AI's potential while safeguarding the core values of integrity, critical thinking, and human intellectual contribution.

We then use the prompt:

Do the Early Career Researchers (ECRs) interviewed here consider Artificial Intelligence (AI) as an opportunity or a threat to their future career?

Early Career Researchers (ECRs) interviewed here generally view Artificial Intelligence (AI) as a double-edged sword, presenting significant opportunities for enhancing their work and career progression, but also posing considerable threats to academic integrity, research quality, and the nature of scholarly work.

Opportunities Identified by ECRs:

Many ECRs are actively exploring and integrating AI tools into their daily work, perceiving them as valuable for:

- Increased Efficiency and Productivity: AI is widely seen as a tool to streamline routine tasks such as drafting and rewriting text, summarising articles, organising ideas, and even preparing grant applications {cited by 11 ECRs} For instance, one ECR found AI invaluable for rapidly producing a high-quality summary of a 600-page book for an urgent project {cited 1}.
- Enhanced Research Capabilities: AI offers benefits in data analysis, pattern identification, and hypothesis generation {cited 8}. Several ECRs highlight its use in their specific fields, such as bioinformatics {cited 1} and cybersecurity {cited 1}, to process large datasets and identify anomalies {cited 2}.
- Overcoming Language Barriers and Improving Writing Quality: For non-native English speakers, AI tools are a major opportunity for proofreading, rephrasing, and translating research, making their work more accessible and publishable {cited 9}.
- Idea Generation and Exploration: ECRs use AI as a "sounding board" or "brainstorming partner" to suggest research topics, explore different angles, and refine research questions {cited 10} (50% of ECRs).
- Reputation Building and Dissemination: AI can assist in improving the visibility of research, designing promotional plans, and tailoring communication for different audiences {cited 8}. Some also see new technological skills (like prompt engineering) as an asset that can enhance their reputation {cited 2}.
- Educational Innovation: AI is being integrated into teaching and assessment to personalise learning, adapt texts for students, and detect errors, potentially reducing teacher burnout {cited 4}.
- Threats and Concerns Expressed by ECRs: Despite the recognised benefits, ECRs articulate significant and pervasive concerns, viewing AI as a substantial threat to their careers and the integrity of science:
- Decline in Research Quality and Critical Thinking: A majority of ECRs (69%) believe that the rapid production of scientific articles facilitated by AI will lead to a decline in overall research quality {cited 1}. They worry about over-reliance on AI leading to a reduction in critical thinking, analytical skills, and intellectual depth {cited 26}. Terms like "watery papers" and "intellectually shallow" are used to describe the potential outcome {cited 3}.
- Academic Integrity Issues: Plagiarism, fabricated data and references, and AI ghostwriting are major concerns {cited 20}. ECRs worry about the difficulty in detecting AI-generated content and the challenge of proving human originality {cited 6}.
- Authorship and Control: Many ECRs affirm that ownership, control, and leadership in research must remain with human authors {cited 19}. They express discomfort with AI being credited as an author, viewing it as a tool rather than a collaborator, and note that existing authorship policies are likely to change to address AI's role {cited 7}.
- Bias and Misinformation: ECRs recognise that AI models, being trained on vast datasets, can embed and even amplify existing biases {cited 21}, E08 (60% of ECRs). There is concern about AI generating confident but false information or "hallucinating" sources {cited 9}.
- Erosion of Fundamental Skills and Job Displacement: Many fear that over-reliance on AI will weaken essential human skills such as writing, logical thinking, and data analysis {cited 20}. Some ECRs express direct fear of job replacement, particularly for entry-level positions or roles involving repetitive tasks {cited 11}.
- Intellectual Property and Data Privacy: Concerns are raised about AI systems using copyrighted material for training without consent or proper attribution, and the potential for research data leakage {cited 18}.

- Devaluation of Publications and Traditional Metrics: A significant number of ECRs (63%) believe AI will change the status or perceived value of publications {cited 1}. This could involve a shift from valuing journal prestige to the utility of content, or a devaluation of scholarship if journals are flooded with AI-polished but superficial work {cited 16}.

In summary, while ECRs, who are predominantly aiming for permanent academic careers [C01, 530], are proactive in adopting AI for its perceived benefits in efficiency and idea generation, they remain acutely aware of its inherent risks. Their responses consistently highlight the critical need for ethical frameworks, transparency, human oversight, rigorous verification of AI outputs, and a renewed emphasis on critical thinking and genuine intellectual contribution to ensure AI serves as a beneficial tool rather than a detrimental force to their future careers and the broader academic landscape.