

## **Generative AI for researchers – a briefing note for the Dean HDR**

Professor Inger Mewburn (contact: [inger.mewburn@anu.edu.au](mailto:inger.mewburn@anu.edu.au))

February 2024

For general circulation

### **Aim**

This paper will highlight the potential uses for ANU researchers of emerging ‘generative AI’ platforms like ChatGPT and co-pilot for Microsoft 365. Even though it is inaccurate, I will use the term ‘GenAI’ as a catch all term for these technologies, which typically involve combinations of machine learning (ML), natural language processing (NLP) and other specialised computer science techniques.

ANU has, for obvious reasons, been most concerned with the use of GenAI in undergraduate teaching and assessment. Researchers have been largely left out of the picture. This paper will focus on what researchers are doing with GenAI right now – and some thoughts on the ethical implications. It is hoped this paper will promote further discussions amongst researchers.

I will concentrate on text for this paper, but obviously there is a wide array of tools for sound and image. If you want a useful guide to currently available technology, this post from Ethan Mollick has [a good summary of the current state of the art](#). It is not the purpose of this paper to discuss GenAI use in Higher Education teaching and learning. There is plenty of material of this sort already; see [this page on the TEQSA website](#).

It should be noted that ANU has a distinguished history of research in the AI field and a lively community of active researchers. This paper is not concerned with what GenAI tools researchers might design and make for their own research purposes.

### **What can researchers use AI for?**

To promote discussion, the table below is a set of use cases. This is not an exhaustive list; it will hopefully serve as a prompt for people to add more. I left out things like drug discovery and modelling and simulation as these are discipline dependent and outside my own field of expertise. I concentrated instead on things any researcher might do with text or data. I included some links to text exchanges with ChatGPT 4 if you want to see how these kinds of text transformation can work in practice.

I left out the ability to write a paper from scratch on the assumption that most researchers are content to keep producing knowledge themselves<sup>1</sup>. I have laid out the various uses in a

---

<sup>1</sup> The introduction of oral exams in July is a useful counter-measure against research candidates ‘faking’ their way through a PhD. I assume that most working researchers are in the business of making knowledge and uninterested in having machines do that work for them, although of course the potential exists for researchers to ‘cheat’ by increasing their research output with GenAI ([see this recent Guardian article](#)). I assume ANU has rigorous enough hiring and HR practices to manage this risk.

table with explanatory notes and thoughts on research ethics implications. I have included more general comments on ethical considerations in my final section.

Use case	Notes	Research ethics implications.
<p>Copy editing and language transformations</p>	<p>Platforms like ChatGPT are very good at performing spell checks, correcting grammar, making text longer or shorter.</p> <p>Writers can also adjust reading level, making text easier (or harder) to understand. Adjusting up makes the text more complex – you can also specify a particular audience and ask the machine to make the text more argumentative. Adjusting the reading level of the text down makes it simpler, good for preparing a presentation to a non-expert audience, for example.</p> <p>The more technical language knowledge the user has, the better the transformations tend to be. Most people will be able to do basic language transformations without training, however my experience is that training helps a lot and have started to incorporate writing with LLMs into Thesis Bootcamp and our productive academic series.</p> <p>There is a growing market of custom GPTs which can have a body of knowledge that informs the answers and structured prompts for producing outputs. These can step users through complex text transformations by asking people to make a series of choices.</p>	<p>Since this use case replicates existing tools like Grammarly, and even language tools in MS Word, this is an uncontroversial use – at least at first glance. If you think about it, there are deeply worrying consequences.</p> <p>We use writing as a thinking tool. As we put words on a page, we are forced to grapple with new thoughts and put ideas into logical and persuasive forms. As GenAI tools become mainstream, people’s basic writing skills are likely to atrophy, in a similar way to how widespread introduction of calculators eroded our ability to mental arithmetic. We no longer need basic mental arithmetic for most tasks, but we do need to use writing to think. The long-term consequences of this are hard to predict. One consequence we are likely to see soon is in the application process. The PhD proposal is a writing (and thinking) test as well as a pitch for the idea – if all our prospective students are all using GenAI it will be harder to assess the abilities they bring into their degree.</p> <p>There is also the potential for algorithmic bias in the way text is cleaned up – so algorithmic transparency and control over the basic assumptions in the text generator is important because both affect the outcome.</p> <p>On the other hand, for people who do not have English as their primary language, or who have dyslexia, the use of these tools could be a powerful tool of equity and inclusion.</p>
<p>Transforming rough notes into ‘camera ready’ text.</p>	<p>Researchers produce notes as they read and think about ideas. Notes are also produced in the field, while doing experiments or while designing and building things.</p> <p>This sort of writing is in the form of ‘notes to self’ that usually no one else sees, and so can be written without academic jargon or proper syntactical structures. Researchers use these notes</p>	<p>It’s easy to do this kind of text transformation on a small set of text, but the toolchain for dealing with large volumes of text is not well developed yet. However, Grammarly already has AI built in and this problem won’t be lying around too long without someone trying to fix it.</p> <p>While this use does not replicate the use of tools like Grammarly and MS Word, the</p>

	<p>to produce more polished writing for supervisors, peers and colleagues as well as presentations. This polishing is a significant time investment for most researchers. A simple prompt such as 'clean up this text for me' will make text good enough to be amalgamated into work in progress.</p>	<p>use is broadly similar and relatively unproblematic. Although</p> <p>The risks of this use are broadly like those noted in the section on copy editing and language transformation.</p>
<p><b>Brainstorming</b></p>	<p>GenAI can be used as a 'sounding board' for researchers looking for feedback on their text as they write. One of the keys to good academic writing is anticipating the readers' objections and questions. If you try to raise and answer these imagined questions and objections in text, it will be more persuasive. Novice writers often forget or don't do enough of this 'reader anticipating,' which is why their text is often unconvincing.</p> <p>You can ask GenAI tools like ChatGPT to help you anticipate reader reactions and look for logical argumentation errors or provide a peer review of the writing. This kind of text transformation is useful if the tool is 'primed' correctly. For example, you can ask ChatGPT to play a role when it critiques your text.</p> <p>For example, I recently wrote a piece on neurodivergence in the PhD and asked ChatGPT to help me anticipate audience reactions of people who might disagree with my argument that PhD examination should change. <a href="#">You can read the text exchange here.</a></p>	<p>This use case is incredibly powerful and will be more so as the technology improves.</p> <p>It can be hard to fact check this kind of output as it is more of a machine generated opinion than a fact. In my testing so far, the machine is at least as good as I am (maybe better) at analysing my text in this way.</p> <p>The critique partner is a role we usually preserve for human supervisors. If the researcher is approaching the machine outputs with their critical faculties attuned, we need to carefully ask ourselves how different this is from asking a peer for their opinion over coffee? An expert peer brings their creativity and (for lack of a better way to put it) 'human intuition' to propositions and problems. Where a machine is helpful for HDR candidates, and working researchers, is to test their ideas before taking it to a broader audience. This is a variation on 'notes to self'.</p> <p>Certainly there is also the risk – again – of being too reliant on machines and possibly missing out on ideas, insights and critiques that are possible with the kind of relational thinking our brains do.</p>
<p><b>Asking for help with a task</b></p>	<p>You can ask ChatGPT to help you diagnose a problem with computer code or how to perform a task in a piece of complicated software like Excel. This is often quicker than googling the problem and watching a you tube video.</p> <p>Researchers can save a lot of time researching background information. Obviously the results need to be fact checked, but ChatGPT can produce helpful summaries on almost any topic you ask for. See <a href="#">this text exchange</a> where I am asking it to tell me about the history</p>	<p>As the technology gets better, this 'ask me anything' function of ChatGPT is likely to get more and more reliable. You can see why Google is slightly freaking out about Open AI taking away its ad revenue from search.</p> <p>Of course, any information produced needs to be fact checked, but researchers have the skills to do this.</p> <p>Simplifying the jargon and insider language is more problematic for three reasons. 1) researchers need to learn to</p>

	<p>of project management for a book I am writing.</p> <p>You can also ask ChatGPT to make complex text simpler. Some disciplines use a lot of jargon and 'insider' language, so ChatGPT can be helpful.</p>	<p>'speak' in the research dialect of their discipline. And 2) in more niche areas one can assume there is not as much text for the machine to learn from and therefore the results will be less accurate and 3) it's harder to fact check an interpretation.</p> <p>On the other hand, for interdisciplinary researchers like me, it is a really fast way to get your head around the language quirks of other disciplines.</p>
Data generation and augmentation	<p>Generative AI can create synthetic datasets that mimic real-world data, helping researchers overcome challenges associated with data scarcity, privacy, and security. This is particularly useful in fields like healthcare, where synthetic patient data can be generated for research without compromising individual privacy.</p>	<p>It's unlikely people would generate synthetic data to analyse, but it can be used to show what the actual data the researchers used might look like</p> <p>The key challenge for researchers working with synthetic data is to verify that it has a close enough resemblance to an authentic set of data to make it useful.</p>
Second language translation	<p>HDR candidates at ANU have reported the use of ChatGPT for language translation and the performance seems to be generally excellent in transformations between, for example, French and English.</p> <p>I only have a (vague) grasp of French and German – both these seem pretty good when I tested passages. I have not tested on less common languages, but I would assume that the smaller the language set in the machine, the more difficult it is to get a good result.</p>	<p>This feature of ChatGPT mirrors the existing technology in Google Translate which is now available as a phone app and in wide use.</p> <p>I don't see the use of ChatGPT to do this as problematic, but obviously how useful it will be varies with language skills.</p>
Help with data analysis	<p>This is an area where we will see an explosion of tools. There is not space here to do this topic justice, I probably need to write a whole other paper.</p> <p>Briefly: the new version of the text coding program MaxQDA has shipped with built in AI for text analysis. Microsoft CoPilot ships in windows now and is available for everyone using that platform. Theoretically this could be used to analyse data in excel, although reports vary as to how good it is at dealing with large data sets.</p>	<p>This is an emerging area and one that challenges research ethics and practice.</p> <p>How much of this kind of work can or should we delegate to machines? Fact checking a data analysis is hard if you can't see how the machine came to the conclusions it presents. There is definitely the danger that researchers just become reliant on these tools and their own analysis skills atrophy.</p> <p>On the other hand, machines do complex things incredibly rapidly and all researchers are under time pressure, so those who do not embrace the technology may well be left behind.</p>

## General ethical considerations for researchers using AI

**Copyright infringement:** Most GenAI companies do not disclose their training sources. Only some (e.g.: Adobe) specifically guarantee that machines have not been trained on copyright material. Copyright law lags behind the technology. The ANU copyright officer currently advises that images produced by AI must be acknowledged, but they are copyright free. This law may change, presenting the most risks for researchers using it in long term research projects (if the law changes before the research is complete, earlier work product may become unusable). The status of ML and copyright has yet to be determined by courts, but there are cases outstanding, the most high profile of which is the [New York Times Vs OpenAI](#)

**Privacy concerns:** For a Large Language Model (LLM) to work, it has to be 'fed' information by the user. If this input is used to train models, there is some evidence this can 'leak' back out (see this New York Times [article](#))

**Hallucination and inaccuracy in output:** Machine learning is based on statistical probabilities with no grounding in truth, bodies of theory or empirical evidence. The maxim 'garbage in, garbage out' also applies. Put simply, Generative AI products cannot always be trusted to produce accurate or appropriate results. The type of training and guard rails built into the models themselves also impact the kind of results researchers can expect.

**Exploited labour:** Most companies that produce AI products are reluctant to reveal 'how the sausage is made', but disturbing reports have emerged about questionable [labour practices](#) in Africa and other places.

**Greenhouse gasses:** ANU is committed to a below zero future, so researcher use of digital infrastructure that is not housed on site is a consideration. The processing required for training ML models happens in large server farms in many different jurisdictions, so the environmental impact is hard to fully assess. However, we do know data centres are massive consumers of energy and their impact ramifies and intensifies as use is scaled up. See for example this analysis of the [digital pathology production chain in the Lancet](#).