AI in the workplace

Six reasons not to use Generative AI tools, and five questions to ask when you use them anyway.

What is this document?

This document came out of discussions in an internal IT working group for a small, tech enabled (but not necessarily tech focussed) company. It arose from the need to come up with some guidance and a practical approach to "AI" as it might impact us in the workplace. It is written for non-specialist "computer users" who are very unlikely to be creating or modifying their own AI tools, but are likely to use or at least encounter them more and more frequently. It was written by Ben Shaw who is probably best classed as a sceptic in this realm, but one that spends a fair amount of time reading about (and playing with) this stuff. It covers:

1. A very brief history and definition[1] of what we're talking about
2. Six reasons why you ought to approach these tools with caution/scepticism
3. 5 principles/questions to ask if you decide to use these tools

This version of the document is open for public consumption, comment, contributions, counterpoints, and general communal use. Please have at it.

Why is this document?

This document argues that:

1. We don't need to worry about an all-powerful AI taking our jobs, enslaving us, or launching all the nukes, but we do need to worry about its real, current impacts on people and the planet.
2. We're talking about something that is neither 'artificial' nor 'intelligent' – these are tools built by and maintained by humans for human purposes. It's important to consider which humans get to benefit, who will bear the costs, and where accountability lies.
3. AI is already baked into so many of the tools we use, and this trend looks likely to continue in the short term at least. If we can't avoid the use of AI, then it's important that we understand how to minimise risk and harm.

What is Generative AI?

"Artificial Intelligence"[2] is a very broad term with a long history, taking in things like image and speech recognition, "big data", chess playing computer programs, and even kitchen appliances. Since the public release of Open AI's ChatGPT 3 in November 2022, most of the discussion (and

[1] This document offers a basic, woefully incomplete, but hopefully practical definition of "AI". For a deeper dive into definitions along the lines of Stafford Beer's "the purpose of a system is what it does", see: https://ali-alkhatib.com/blog/defining-ai

[2] There's a lot of scare quotes in this document, because these terms are heavily contested and often based more on marketing hype than sober definition. Suffice to say there's plenty of reasons to believe that Generative Artificial Intelligence is not Generative, Artificial, or Intelligent, some of which are covered in this document!

this document) has focussed on what is called “Generative AI”. Generative AI (hereafter genAI) is the name given to a suite of tools into which you can put some kind of “prompt” and receive in return a text, visual or audio output, that resembles something that a human might produce. They differ from more traditional search technologies (which also might employ various forms of “AI”), in that the outputs they create are “novel”[3] – that is they are created on the fly in response to the prompt given, not just pulled from storage based on some matching criteria. Some well-known examples of genAI in use at the moment are OpenAI’s ChatGPT and DALL-E, Midjourney (Midjourney, Inc) and Microsoft’s Co-pilot (built on ChatGPT). If you want to nerd out a bit on how at least Large Language Models (LLMs) like ChatGPT work, I can recommend this article. The short version is “spicy autocorrect”.

Whether or not we *want* to use genAI, the monopolies that Microsoft, Google, Meta, Adobe, Amazon and other large software providers have, means we’re unlikely to have much choice in using at least some forms of this technology. It is already, or soon will be, baked into our phones, computer operating systems, word and data processing tools (not to mention our cars, our homes etc). This is not a reason to throw up our hands and accept the inevitable, but a call to take extra care with how and where we use these tools, and perhaps to join the growing chorus of voices arguing for a much broader and richer conversation about how we want this technology to impact our lives.[4]

Six reasons why we should hesitate before using genAI

A bunch of great articles have been written about the risks and dangers of AI (and the sometimes uh, “fringe” beliefs of its chief proponents), that I won’t try to summarise here (there’s some links at the bottom!). Suffice to say these are real, present, negative impacts that are happening in the world right now, not theoretical ones about god-computers.

1. **genAI amplifies and feeds biases and misinformation.** It was raised on a giant bucket of the messy human stew we call “the internet”, and there are countless examples of the ways this plays out in what it generates. This is not just a problem of needing “better data” (in short, we need better humans, but that’s a longer-term project). In addition to the more blatant sexism and racism, it favours particular forms of language (mostly English) over others, creating a very real risk of homogenisation of the “correct” way to speak and write.
2. **Most well-known genAI models are built on a bunch of stolen data**, (and could still be stealing yours). It’s still not clear who will be holding the can when the inevitable lawsuits start arriving in force. Maybe you?
3. **It's designed to appear to be “helpful”, not to be accurate**. It doesn’t so much “hallucinate” as just give you an answer it considers most “likely”. This will usually (not always!) be right in simple questions like “what is 2 + 2” (maybe), but not so much in questions like “what crimes have KPMG committed lately?” or “can you melt an egg?”.
4. While the term “artificial intelligence” conjures images of machines independently consuming and producing data, **it actually obscures an enormous amount of poorly paid and exploitative human labour.** Someone’s gotta train the model, tell the machine what’s “objectionable”, and sometimes even pretend to be the computer. The need for speed, scale and exposure to traumatic material required means this work is being done for very little money in places with very few labour protections.

[3] Though exactly how novel is up for debate, a better term might be “statistically likely”

[4] It’s worth pointing out that that nothing in technology is really inevitable, remember when we were all going to have 3d printers/3d TVs/the last version of AI/cryptocurrencies/NFTs/virtual reality goggles etc etc etc).

5. **AI processes consume vast amounts of electricity, water and material.** According to an estimate in July 2023, the daily energy use of ChatGPT alone is equivalent to 33,000 US households.
6. **It might yet collapse/become very expensive/turn into something terrible**. There are few things more precarious than the next "trillion dollar idea" that doesn't actually have a business case. Right now the cost of running these models is astronomical (and growing) so the free ride is likely to end soon. Whether that happens before or after the user base collapses or the models consume themselves is another story. A lot of very big players are very exposed, and OpenAI in particular isn't doing so great at time of writing[5].

Five questions to ask when using genAI

So given that there's some air of inevitability (at least in the short term) to all this, and that some portion of people will have skipped straight to this bit - what do you need to keep in mind when using Generative AI tools?

1. **Do I even need to write this?** If you're using a genAI to automate some piece of tedious, pro-forma text, perhaps ask yourself, "does this text need to exist at all?"[6]. If you don't care enough/have time/expertise to write it, will anyone care enough to read it?
2. **How do I know that it's true/accurate?** Fact check anything that you care about the answer to. I would go as far as to suggest "nothing goes to a client without a human confirming its veracity"[7]. genAIs are designed to generate likely outputs, not facts. And no, asking the machine to fact check itself doesn't work.
3. **Am I giving up proprietary and/or private data?** Don't enter or produce anything that you want to maintain any kind of ownership over. Firstly, courts are unlikely to award IP rights to content produced by genAI. Secondly, while most companies have given vague assurances that they won't steal your data (particularly on paid plans), their track record isn't great in this area, and there's always a chance that what you thought was private was actually feeding a very public interface. Even Microsoft who owns a big chunk of OpenAI is worried about this.
4. **Am I being transparent about using genAI?** Always specify any content that was generated (or assisted) by a genAI. This is basic courtesy to anyone reading the content who may want to treat it with an additional level of scepticism, but also the terms and conditions of some tools explicitly prohibit you from passing off genAI content as human created.
5. **Am I building a house on sand?** Don't build any critical infrastructure or processes on anything that is currently free. However you feel about the chances of the whole system collapsing in the near future, it is going to need to make (a lot of) money at some point. Recent history gives us reason to think that we the users are most likely to lose out in any monetisation scheme, but in simple terms you should ask "if this cost $10 a pop to do, could I still do it? What about $100?".

---

[5] Any link I include will be obsolete by the time you read this. Go to your search engine of choice or ask chatGPT..
[6] Noting the point above about energy - every bit (and byte) of data we create has an ongoing carbon cost
[7] As one example, PWC, who are of course happy to push genAI to their clients, have banned its use for client work

Further Reading

- This excellent list of articles curated by @poisonivy47.bsky.social pulls together writing/research on just about every aspect of genAI and its impacts
- Per Axbom's Elements of AI Ethics is a useful unpacking of the various dimensions of harm occurring right now, that need to be considered when thinking about "ethical" AI
- This interview with Emily Bender does a really good job of covering some of the basics of what genAI is, in plain language
- A deep, but still surprisingly accessible intro to how Large Language Models generate "novel" text from Stephen Wolfram
- This paper from the Norwegian Consumer Council on potential consumer harms of AI is a great overview of everything talked about in this document, and more.