AI companies must own up to their sources

'Nothing will come of nothing,' King Lear said. For bots, that means mining centuries of human endeavour and literature

Hugo Rifkind, Tuesday March 12th 2024, The Times

The rise of artificial intelligence — of computers that appear to think and create — is fascinating, frightening, existentially challenging and a bunch of other things, too. In the end, though, there is one fundamental fact that we must always remember. Which is that what comes out is based on what goes in. Or, as King Lear put it to Cordelia, "Nothing will come of nothing."

An example. Last month, OpenAI unveiled Sora, a model that generates video from text. The first samples released were startling. The company published both the prompts that it had fed into its new system — show an old man thinking in a Paris café, show California during the gold rush, show a cartoon otter surfing, etc — and the results, which looked like trailers for Hollywood films. When you realised what you were looking at, you could only gulp. It was one of those moments where you suddenly have a flash of being balanced not only on a technological precipice but on a human one too.

Before long, though, an oddity emerged. A bunch of AI enthusiasts — I think a creative director called Nick St Pierre was the first — had taken the text prompts given to Sora and fed them into AI generators made by rival companies instead. These were now still images rather than video but the results were uncannily similar. As in, much the same old man, much the same Parisian café. A near twin adorable cartoon otter. Leading many to wonder, why?

Remember, nothing will come of nothing. Remember also that what comes out is based on what went in. Here, that obviously included the prompts, which in fairness were more detailed than I've outlined above. Perhaps, once you've specified that the old man wears a brown beret, that the lighting is cinematic and so on, then different human artists would produce similar results too. But there is also the question of background inspiration. What databases of past works did these models have at their disposal? What films, what photos, what novels, what paintings? Might it, each time, have been the same one?

The Times has reported on a fight between Britain's biggest publishers and the tech industry. According to the Publishers Association, members including Random House, HarperCollins and Oxford University Press believe that "vast amounts of copyright-protected works" are being fed into text-generating AIs to create similar databases without authorisation. And, more vitally, without recompense. Almost certainly, they're right.

This is one of a host of similar fights being fought by creative industries. Their case makes sense. Base a film on a book, after all, and you pay. Copy music, likewise. With AI, it should be possible to unravel the knot and see what has come from where. So, if it could be proven that a new AI-generated film took even a fraction of a per cent of its ingredients from, say, the Indiana Jones franchise, then George Lucas deserves a cheque. Likewise with all the other fractions from everywhere else.

Yet what complicates matters, headscratchingly, is that humans don't create in a vacuum either. A singer may be inspired by Bob Dylan, or a children's author by JK Rowling. Is that so different? Even for this article, indeed, stuff went in before stuff came out. In my life to date, I have read, seen, sought and stumbled across the works of countless other people, and my own thoughts would be different otherwise. So should I be giving part of my salary to them?

With my brain, or yours, these considerations swiftly become ineffable. Are we meat machines, differing from AI only in our complexity? Maybe. That's one for the philosophers. With AI, though, the inputs are only ineffable if the tech firms want them to be. And, right now, they do.

This is not just about art. Last month, executives from leading news publishers — including News UK, which publishes The Times — appeared before the House of Lords communications and digital committee, pleading for stronger safeguards against AI devouring and then replicating news content, again without permission or control.

Even copyright, though, is just the tip of the AI iceberg. In November The Lancet reported on a new AI tool, developed in London, that identifies cancerous growths from CT scans more effectively than any humans can. Obviously, I'm not suggesting this is bad news. Again, though, the AI is not doing this in a vacuum. Sucked into the system will be untold hours of past human labour and expertise. Generations of doctors, peering at scans. Textbooks, lectures, classes and experiments, stretching back into the dark. Don't forget them.

The same is true in almost any educated field you can name. In law, in accounting, in engineering, in psychology, in coding, governance; in whatever, we're on the threshold of a time that, whatever you do, an app will do quicker and better. Don't kid yourself otherwise. For humanity this is a boon, much as it might feel otherwise if this is currently how you earn your keep. By the time the bots are better at your job than you are, it may be too late to ask how they got that way.

What we need to remember, and what tech firms need to remember too, is that these super, thrilling, boggling machines are only ever the very last part of the story. Is a computer inspired by consuming literally everything less creative and more parasitical than a human inspired by only consuming a bit? Like I said, that's one for the philosophers. The difference now is that we ought to be able to see it happening, and the only reason we cannot is that AI firms don't want us to.

So we must make them. Make them keep the receipts, and make them publish them, and insist they make good for what they have taken. Remember, nothing will come of nothing. And if that means I owe Shakespeare a fiver, well, so be it. Send him my way.