

The ghost in our machine: How AI is quietly changing us

Aryan Chopra

The AI era seems to have arrived without a bang. One moment, it was just something for fantasy, and the next, it was subtly integrating into our lives. It is the invisible power that is indicating a strange charge on your credit card, identifying a minute shadow in a medical scan, or recommending the next song which you would play repeatedly for a week. It writes even the uninteresting reports that we try to avoid. They are all very helpful prompts but sometimes, they can be so drastic that they change the whole situation. And occasionally, they misinterpret the situation completely. The question now is not whether we should use AI, the answer to this question is already clear. The more relevant question is, how do we coexist with it without the loss of the human traits that define us?

AI is not a disruption in the world of politics that is still being decided by some far-off debate. It is changing the way we think and act when using the Internet, one click at a time. It is a silent but powerful rewiring of three main things that we often overlook: our attention, memory, and choices. First, our attention. It's being sliced and diced. The feeds and recommendation engines on our screens are designed for one thing: to keep us watching. They learn what makes our eyes stick and then feed us an endless buffet of it. Great for their business model. Terrible for our ability to focus. Our attention spans get shorter, and we end up in echo chambers that make public conversations feel more like shouting matches. Then there's our memory. You've probably heard of the "Google effect." It's that feeling of knowing you don't need to remember a fact, just where to find it. We're offloading our brains to search engines and digital assistants. It's efficient, sure, but it comes with a risk. If we stop doing the hard work of remembering, we might forget how to build real expertise in the first place. It's like we've all agreed to outsource a part of our own minds. Finally, we're letting machines make our choices. Algorithms now have a say in who gets hired, who gets a loan, and even who gets priority in a hospital. These systems are powerful because they learn from history, but that's also their biggest flaw. If the historical data is full of biases, and it always is, then AI will also learn to be biased, too, but with terrifying efficiency. A decision can feel perfectly logical and be deeply unfair at the same time.

These three shifts are changing us long before we even start talking about AI regulations. We notice the convenience, but we don't always see the slow erosion of our own agency. At the office, AI is a productivity machine. No doubt about it. Studies have shown that tools like GitHub's Copilot help developers blast through coding tasks way faster than before. The same goes for anyone who must write reports or answer customer emails. AI can handle the first draft, freeing us up for other things.

But "productivity" isn't just about speed. The interesting thing is, it seems like junior employees get the biggest boost, which helps close some skill gaps. The danger, though, is that they might start using AI as a crutch. If you never have to do the boring, routine stuff, how do you develop the deep, nuanced judgment that only comes from experience? When the machine does all the routine thinking, we lose the chance to learn. Then there's the creepy side of it companies using AI to survey workers or just replacing them to cut costs. It leaves everyone feeling insecure. So, the future of work isn't just about getting more done. It's about completely rethinking what our jobs are. The companies that use AI to help their people think better will probably do well. The ones that just see it to cut headcount. They might find they've hollowed out their own expertise. It's easy to see why schools are jumping on the AI bandwagon. Adaptive learning tools can give every student a personalized exercise and instant feedback, freeing up teachers from the drudgery of grading. In a classroom of thirty kids, that's a huge deal.

However, there is a downside to it. The healthy struggle that aids in learning gets missed by students if they get into the habit of letting an AI think for them. It cannot be simply banned; that would be a battle one would lose. The solution lies in educating students to use it as a tool, not as a competitor for their own intellect. We still need to make them solve problems the old-fashioned way. We must

hammer home the importance of checking sources and building a real argument. And we must make sure it's fair. If these AI tutors are trained on data from only one group of students, they could end up giving bad advice to everyone else. Access must be for everyone.

Nowhere is AI's promise and peril more obvious than in healthcare. In some tests, AI models can spot tumours on scans as well as a human specialist. They can help hospitals figure out who needs a bed most urgently. But the devil is in the details. There was a famous case where a widely used algorithm was found to be systematically giving Black patients less care. Why? Because it was using how much a person had spent on healthcare in the past as a stand-in for how sick they were. Since Black patients historically had less access and spent less, the algorithm thought they were healthier than they were. It was a perfect, logical system that ended up amplifying a massive structural inequality, affecting millions before anyone caught it. That story is a huge wake-up call: the goals we give an AI matter just as much as the code it runs on. Regulators are trying to catch up, but the tech is moving so fast that the rules are always a few steps behind. We need human doctors in the loop, constantly checking the machine's work.

AI is doing some good for the planet. It is aiding us in the forecasting of weather patterns, the enhancement of power grids, and the monitoring of deforestation through satellite images. Yet it has a hidden price: training the giant AI models consumes mind-boggling energy. Some say that the power requirement for data centres might overtake AI supply by 2030. It's like AI is not just an application anymore. Rather, it has become a physical entity with an environmental footprint. Certainly, we can improve the situation by implementing advanced cooling technologies, however, we must start incorporating sustainability into AI development as a core value rather than as an afterthought. Generative AI has completely disrupted the creative industry. With the help of AI, anybody can generate images, music, and text. It is a wonderful tool for brainstorming or for a beginner to get the initial push. But the legal and ethical issues are just beginning to unfold. Is it justifiable to use an AI model which has trained on billions of images from the internet without asking the artists for their consent? Multiple lawsuits are leading the way. And on the cultural front, when algorithms start to choose the art and music that we see, they are also moulding our taste in a very subtle way.

A few simple norms could help. Be transparent about when AI was used. Figure out a way to compensate artists whose work was used for training the AI model. Build platforms that show where an image or song came from. This isn't just a tech debate; it's a conversation about what we value in human creativity. Governments are scrambling to figure out how to regulate all of this. The EU has its AI Act, which puts strict rules on high-risk systems. Other places are trying a lighter touch. This patchwork of laws means companies can just shop around for the weakest rules, leaving a lot of people unprotected. Global standards, like UNESCO's recommendation on AI ethics, are also emerging to guide this process.

Good rules should be cantered on people. For the high-stakes stuff things that affect our lives, our freedom, our jobs we need tough, independent testing and clear explanations for how these systems work. And when they get it wrong, there must be a way for a real person to step in and fix it. For everything else, transparency is still key. The more we can involve everyday people in these conversations, the more trust we'll build.

We can't just unplug it all. But we can be smarter about how we live with it. Here are a few ideas:

1. Ask what the machine is really trying to do. Is it optimizing costs? Or engagement? Or for a real human benefit? Bad goals make bad AI.
2. Protect your unplugged brain. We need to carve out spaces in our schools and workplaces for thinking and creating without AI's help. That's how we build the ability to make judgments.
3. Keep checking its work. AI models aren't "set it and forget it." They need to be constantly monitored for bias and weird side effects.

4. Demand a human in the loop. There must be a way to appeal an automated decision to a real person.
5. Think about the environmental cost. Ask about the energy footprint of the AI services you use.
6. Teach everyone how it works. We all need a basic understanding of what an AI is doing so we know when to be sceptical of its advice.
7. These aren't magic solutions. But they shift the conversation from "the AI is coming for us" to "we get to decide how AI serves us."

AI mirrors the human being. It is boosting the best traits humans have and, at the same time, it is exposing the worst ones. It can provide accurate disease diagnoses and suggest the best educational practices, but on the other hand, it can also give a boost to our prejudices and scatter our attention. The main factor is not the technology that we have; it is us. It is the decisions we make regarding the purpose of these tools, the people involved in their development, and the limitations we set in terms of what we are going to safeguard. Coexisting with smart machines does not mean putting up with their dictates. It is our duty to create the systems that will enable us to become more powerful, more just, and more humane. We are the very first generation that has such a close contact with AI. This could be our very first and the only opportunity.