

The Intelligence Age by Sam Altman

September 23, 2024

In the next couple of decades, we will be able to do things that would have seemed like magic to our grandparents.

This phenomenon is not new, but it will be newly accelerated. People have become dramatically more capable over time; we can already accomplish things now that our predecessors would have believed to be impossible.

We are more capable not because of genetic change, but because we benefit from the infrastructure of society being way smarter and more capable than any one of us; in an important sense, society itself is a form of advanced intelligence. Our grandparents – and the generations that came before them – built and achieved great things. They contributed to the scaffolding of human progress that we all benefit from. AI will give people tools to solve hard problems and help us add new struts to that scaffolding that we couldn't have figured out on our own. The story of progress will continue, and our children will be able to do things we can't.

It won't happen all at once, but we'll soon be able to work with AI that helps us accomplish much more than we ever could without AI; eventually we can each have a personal AI team, full of virtual experts in different areas, working together to create almost anything we can imagine. Our children will have virtual tutors who can provide personalized instruction in any subject, in any language, and at whatever pace they need. We can imagine similar ideas for better healthcare, the ability to create any kind of software someone can imagine, and much more.

With these new abilities, we can have shared prosperity to a degree that seems unimaginable today; in the future, everyone's lives can be better than anyone's life is now. Prosperity alone doesn't necessarily make people happy – there are plenty of miserable rich people – but it would meaningfully improve the lives of people around the world.

Here is one narrow way to look at human history: after thousands of years of compounding scientific discovery and technological progress, we have figured out how to melt sand, add some impurities, arrange it with astonishing precision at extraordinarily tiny scale into computer chips, run energy through it, and end up with systems capable of creating increasingly capable artificial intelligence.

This may turn out to be the most consequential fact about all of history so far. It is possible that we will have superintelligence in a few thousand days (!); it may take longer, but I'm confident we'll get there.

How did we get to the doorstep of the next leap in prosperity?

In three words: deep learning worked.

In 15 words: deep learning worked, got predictably better with scale, and we dedicated increasing resources to it.

That's really it; humanity discovered an algorithm that could really, truly learn any distribution of data (or really, the underlying "rules" that produce any distribution of data). To a shocking degree of precision, the more compute and data available, the better it gets at helping people solve hard problems. I find that no matter how much time I spend thinking about this, I can never really internalize how consequential it is.

There are a lot of details we still have to figure out, but it's a mistake to get distracted by any particular challenge. Deep learning works, and we will solve the remaining problems. We can say a lot of things about what may happen next, but the main one is that AI is going to get better with scale, and that will lead to meaningful improvements to the lives of people around the world.

AI models will soon serve as autonomous personal assistants who carry out specific tasks on our behalf like coordinating medical care on your behalf. At some point further down the road, AI systems are going to get so good that they help us make better next-generation systems and make scientific progress across the board.

Technology brought us from the Stone Age to the Agricultural Age and then to the Industrial Age. From here, the path to the Intelligence Age

is paved with compute, energy, and human will.

If we want to put AI into the hands of as many people as possible, we need to drive down the cost of compute and make it abundant (which requires lots of energy and chips). If we don't build enough infrastructure, AI will be a very limited resource that wars get fought over and that becomes mostly a tool for rich people.

We need to act wisely but with conviction. The dawn of the Intelligence Age is a momentous development with very complex and extremely high-stakes challenges. It will not be an entirely positive story, but the upside is so tremendous that we owe it to ourselves, and the future, to figure out how to navigate the risks in front of us.

I believe the future is going to be so bright that no one can do it justice by trying to write about it now; a defining characteristic of the Intelligence Age will be massive prosperity.

Although it will happen incrementally, astounding triumphs – fixing the climate, establishing a space colony, and the discovery of all of physics – will eventually become commonplace. With nearly-limitless intelligence and abundant energy – the ability to generate great ideas, and the ability to make them happen – we can do quite a lot.

As we have seen with other technologies, there will also be downsides, and we need to start working now to maximize AI's benefits while minimizing its harms. As one example, we expect that this technology can cause a significant change in labor markets (good and bad) in the coming years, but most jobs will change more slowly than most people think, and I have no fear that we'll run out of things to do (even if they don't look like "real jobs" to us today). People have an innate desire to create and to be useful to each other, and AI will allow us to amplify our own abilities like never before. As a society, we will be back in an expanding world, and we can again focus on playing positive-sum games.

Many of the jobs we do today would have looked like trifling wastes of time to people a few hundred years ago, but nobody is looking back at the past, wishing they were a lamplighter. If a lamplighter could see the world today, he would think the prosperity all around him was

unimaginable. And if we could fast-forward a hundred years from today, the prosperity all around us would feel just as unimaginable.