

THE 2024 AI DIRECTION OF TRAVEL

PROFESSOR ESWAR PRASAD

Professor, Cornell University

November 14, 2023





AI has generated a high level of excitement, uncertainty, fear, and even speculative frenzy. These are typical of the opening act of practically all transformative technologies. Before we assign that label to AI, however, it is worth considering what frames of reference we should use to evaluate the promise and potential of a new technology.

One approach is to ask if a technology transforms things in a fundamental way including creating new products, new ways of conducting old forms of business, and changing the way humans interact with each other and with the technology itself. Or does it simply provide a more efficient way of conducting existing forms of business and of making current products and

services better.

Since we're at the Singapore FinTech Festival, perhaps it's worth comparing what we see in the world of FinTech and especially the cryptocurrency revolution, and line it up with AI. The parallel is far from exact because cryptocurrencies were meant to merely transform the world of finance. But it is worth thinking about whether they really had a transformative effect. It is an open question what the legacy of the cryptocurrency revolution will be, whether what it really did was identify a set of deficiencies in existing financial systems and serve to catalyze changes. That is transformative but not in the sense of the technology itself being transformative in its own right.

“

AI certainly goes far beyond that even in its ambition. It is about changes in financial services. It's about how healthcare can be conducted in better ways, using big data to improve patient outcomes. It's about how education, my own field, can be supercharged in ways that benefit students but even the very process of knowledge creation.

”

AI is thus certainly going to be much more pervasive. In terms of principles of evaluating new technologies, though, perhaps it is not that different.

“

The promise of AI is undeniable and the technology certainly has the potential to be transformative. But at a time when it is tempting to get carried away by the promise, it is worth considering some possible deficiencies as well.

”

Let us then consider a few paradoxes where I will draw symmetries between the cryptocurrency revolution and the AI revolution, and then talk about what role governments might play in terms of creating guardrails to harness the benefits of this technology.

The specific nature of AI becomes important as we consider exactly what sort of standards we should employ in evaluating this new technology. For predictive AI, some of the standards are fairly clear. One can think about accuracy, perhaps reproducibility of certain results. If you think about using

AI, in, consumer service settings, the accuracy and latency (time delay) of the responses that you get from an AI are important benchmarks. This is a more or less well defined set of criteria, where one can build specific quantitative benchmarks.

Things become a little more complicated when we think about generative AI, because now we're talking about new content, new products, new services rather than just improved versions of existing ones. The element of creativity is inherently much harder to evaluate and judge. There is the additional question about whether there is a way of scaling this output and developing quantitative metrics for determining how these new products and services or inventions can improve welfare for consumers, businesses, and perhaps even governments.

If you think about the health care field, certainly finding a new cure, or a new way of combining certain chemicals in order to create a new chemical compound that has curative or other desirable properties, that certainly is something that generates well defined benefits. What about an AI creating a new piece of art, a new piece of music. Those are remarkable creations in their own right. But what is the correct standard now? When we think

about recreating a particular picture, reproducibility is a clear criterion. But when you're creating something completely new, that standard is no longer relevant.

A much harder question is whether there is anything intrinsic in terms of generative AI that can match human values. Many technologies are inherently value free. But when we think about self-generating mechanisms or corrective learning mechanisms built into AI, we're asking them to essentially evaluate themselves against value structures that we humans have embedded in our psyches. Those values are typically not codified in any form that an algorithm can learn from observed data. This strikes me as a potential problem. If we have technologies that are let out into the world, that really can transform and create new things. But there is no way in which one can think about the values or value systems embedded in those new creations, that is a concern. Here again, there are some parallels with the world of cryptocurrencies.

Some of these parallels extend to even more basic considerations about which directions new technologies will take us in. Blockchain technology is marvelous but blockchains are self-contained ecosystems that don't communicate very well with other blockchains or

with the outside world. This creates a need for oracles, pieces of software that convey information between blockchains or between a particular blockchain and the real world. Oracles, it turns out, are an important point of vulnerability given their current state of development. Thus, we take the very nice and secure world of a blockchain and, in the process of increasing its functionality, expose it to external vulnerabilities. There is an analogy here when we think about generative or predictive AI.



The sort of data that are fed into an AI can determine what you get out of it. That intrinsically creates a filter, where the right sort of data can give you the right sort of answers, but therein are already embedded certain value judgments about what sort of data you use, who provides the data, how those data are used, and so forth.



As we think about other elements of both predictive and generative AI, there are concerns about whether these constitute a substitute for human intelligence in some form. In my own field, in education, we are beginning to see ChatGPT play an important role. I don't forbid my students from using ChatGPT for writing their papers. And certainly they seem to be writing much clearer papers with fewer errors. What I worry about is whether this reliance on technology affects their critical thinking. One could argue that human beings can do the critical thinking and then leave the unexciting stuff, you know, the writing-up and polishing of those ideas to AI. What I have found in my own experience is that when I'm standing before you, thinking about how I convey these ideas to you, that's what forms my ideas and sharpens them. It's when I endeavor to write down certain ideas that those ideas take shape and become clearer. So now when my students think they can use ChatGPT but cordon that off from the critical thinking, I worry that that doesn't happen because it is the process of writing, it is the process of articulation of ideas in one's mind or in public, that really forms those ideas. There are some dangerous issues here as we think about where this will lead us. The ability of critical evaluation might be lost if we stop using our critical thinking and, instead, surrender some of that ChatGPT.



This is not to say that AI doesn't have a role, but it bears careful thinking what path it is taking us down. I alluded earlier to some paradoxes and how there is an interesting parallel with the cryptocurrency world, which is what we would have been talking just like last year if ChatGPT had not emerged on the stage in the last few months.

One of the issues is related to whether we have actual decentralization thanks to the cryptocurrency revolution, which was the whole point of decentralized finance. One could make the case that decentralization of finance and broadly accessible technologies like AI can level the playing field among different companies. One could

imagine startups being able to scale up their business models easily using these technologies, which even give individuals and low income countries the ability to start competing toe to toe with more established players. The reality, though, is that we're witnessing a lot more centralization. Much of what works well in the cryptocurrency realm is actually quite centralized. After all, stablecoins, centralized exchanges and other such mechanisms take us away from rather than toward a decentralized architecture. One could envision a parallel in the AI world, where a few companies such as Open AI dominate the field. So again, rather than decentralization and more competition, you could end up with much more concentration.

Another paradox is related to trust. Bitcoin and the cryptocurrency revolution were intended to reduce our reliance on institutionalized trust. What we've learned is that humans do want things to trust, which is why we have people, although they could keep their money in decentralized wallets, essentially keeping all their money in exchanges that they can see and that have a name to them, like FTX, Binance and so forth. We see much more centralization of various forms of trust when it comes to trading cryptocurrencies.



ChatGPT and AI more generally have the potential of changing trust, or one might argue that they can actually build trust. But one can equally well imagine a world where they end up making it much harder to trust anything. They could actually end up undermining trust, and making it much harder for societies to function in a reliable way.



Yet another paradox is that many of these technologies are meant to reduce the power of governments, traditional financial institutions, and big corporations. That's not what we're seeing in the world of finance. With decentralized cryptocurrencies such as Bitcoin not working well as mediums of exchange, we are moving to a world wherein central banks are issuing their own digital currencies—retail central bank digital currencies. Then we have stablecoins issued by specific companies, with trust in their stable value coming from their reliance on fiat currency backing. These developments could result in large corporations and governments having even more pervasive influences in our economic and possibly even social lives. Could AI take us down a similar road, making it much harder for us to escape the intrusion in our lives from governments, corporations, and other official and private institutions?

The parallels between the cryptocurrency and AI revolutions in terms of these paradoxes should not be oversold. They do point to one crucial lesson, though. As we look at different technologies, the proposition that technology can police itself, can take care of the problems it creates if you leave it to itself, is not a tenable one. This brings up the role of governments and what role they are going to play. We've heard innovators

and entrepreneurs in AI saying they want more regulation. They seem to be clamoring for regulatory clarity. We've seen this time and time again, what regulatory clarity often seems to mean to innovators and entrepreneurs is the conveyance of legitimacy to the new technology, with relatively light touch regulation so that that technology can prosper relatively unfettered from actual regulation. This is clearly not an ideal outcome.

For open-minded regulators, those who view these technologies as potentially having a useful purpose in society and want to create guardrails, the Do No Harm approach is often touted as critical to creating space for innovation. New technologies inherently come with risks. So the notion of Do No Harm ought to be replaced with manage the harm or contain the harm, because there is going to be some harm. The question is whether that harm has systemic consequences, or if the consequences of those risks fall on those who are least able to absorb those risks. If you think about the cryptocurrency revolution, investor protection was very important because you could have a lot of naive, unsophisticated retail investors taking on risks that they didn't understand and, in some cases, couldn't manage the consequences of. If it was more

sophisticated investors, perhaps venture capitalists, taking on such risks that's not a problem unless there are systemic consequences. So the issue is how to balance regulation to not avoid risks but create sufficient protection for certain classes of investors.

Another reality is that we're going to have to think about governance issues not at the national level but at the global level. This again raises a whole host of questions. We've had a lot of progress in terms of specific jurisdictions moving forward in AI regulation. The EU, in particular, seems to have gotten out ahead of the rest of the field in this area. The Biden administration has just issued an executive order on safe, secure, and trustworthy AI. A number of other jurisdictions are also beginning to undertake moves on the regulatory front.

The reality, though, is that we're going to need a global governance architecture for governance of this new technology to be effective. With that, as with any issues related to global governance, brings up a host of complications--who gets to write the rules, who gets to be at the table when the rules are written, and who gets to enforce the rules? These are issues that are complicated in every aspect of global governance. AI is no exception to that rule but its emergence at a time



when geopolitical tensions are running high will make this a fraught process. The challenge for regulators worldwide is to develop a cohesive regulatory structure rather than one that is fragmented.

One of the issues about AI is that it is not just value free but also free, at least in principle, of the sort of pettiness, the prejudices, the tribalism, that characterize us as human beings, especially when we approach issues at a cross national level.



So, on a more positive note, as we think about global governance problems that look intractable, maybe what we really need is the hand of AI to move us toward a better state of the world.



That would be a great contribution, indeed, from a technology that was devised by human beings but might allow us to rise beyond the differences that hold us back.