



**“Artificial
intelligence (AI)
isn’t just poised
to change the world;
it’s already changing
the world”**

Technology and the Future

Haroon Sheikh

Haroon Sheikh is professor of philosophy at Vrije Universiteit Amsterdam, where he teaches on globalization, the philosophy of technology, and East-West geopolitics. His books address a diverse array of subjects, including relations between Europe and Asia and the politics of water. Alongside his academic career, Haroon advises the Dutch government as part of the Scientific Council for Government Policy (WRR).

Artificial intelligence (AI) isn't just poised to change the world; it's already changing the world. Across a wide range of sectors, powerful AI technologies are finally leaving the lab, proving their commercial viability, and transforming entire industries. For some, this change is a source of alarm, but I want to add some nuance to the conversation about AI and provide a balanced, historically informed perspective of how AI technologies are likely to develop in the coming years. At the outset, it's important to draw a distinction between artificial general intelligence (AGI) and narrow AI. The former



is the attempt to recreate the human mind, which has been a dream of computer scientists since the earliest days of the AI discipline. Progress on this front remains slow, mostly because our understanding of our own consciousness is still so limited. It's incredibly difficult to recreate something that you don't fully understand in the first place, particularly a system as complex as the human mind! Given the scale of the challenge, I don't expect generalized intelligence to emerge any time soon – but AGI is only one part of the bigger picture. There's also the larger and more promising field of narrow AI – that is, an AI that can perform a particular function better than a human. This arena is where the most exciting developments are currently taking place.

Within narrow AI, we're seeing the emergence of tools that don't just interpret patterns, but create new ones – a huge step forward. One such development is ChatGPT, an AI-powered chatbot developed by Open AI that can provide lengthy, detailed, and considered written responses to user prompts. The same goes for DALL-E, Open AI's image generation tool: it creates works of art based on a written prompt, and the results are truly outstanding. Tools like these are set to become mainstream within the next few years.

AI also has huge potential when it comes to federated learning. Perhaps most interestingly, I think this technology addresses many of the privacy concerns that surface in the AI debate. In federated learning, instead of feeding large quantities of data to a central algorithm, the algorithm migrates to different data sources. For instance, the AI could visit the server of a particular hospital, adjust its parameters based on what it learns from the data, and then return without taking any data with it. As such, I think federated learning will contribute to the wider adoption of AI technologies in sectors – like the healthcare industry – that are more skeptical toward data-driven approaches.

The story isn't uniformly positive, though. Other developments in narrow AI are more problematic, such as the rise of generative adversarial networks – the technology behind



Artificial intelligence isn't just poised to change the world; **it's already changing the world**

deepfakes. The newest deepfake tools allow users to create fictitious videos with remarkable accuracy, and this has serious political consequences. We already live in an era of fake news and filter bubbles, and in a world of deepfake content, the line between facts and “alternative facts” will become even more blurred.

However, I don't think this is inevitable. With the proper regulatory approach, governments can find a path that restores public trust in AI and addresses its potential downsides. A good example of this is algorithm registration, which is already required in cities including Amsterdam and Rotterdam. Under this system, government organizations must publish all the algorithms they use. Let's say a citizen has received a fine from the municipality: with algorithm registration, the citizen will be able to see whether the decision to fine them was made by a human or an algorithm. If it was an algorithm, they'll be able to see which rule was being applied. I think public confidence in AI will improve substantially once citizens are no longer in the dark about what rules are being applied and how.

I'm also encouraged by the fact that EU policymakers are reaching a consensus about AI. After seemingly endless

declarations, concepts, and roadmaps, the EU AI Act is set to become binding law within the next few years. This first-of-its-kind legislation will ban applications and systems that create an unacceptable risk, place specific legal requirements on high-risk applications, and allow low-risk technologies to operate with minimal regulation. Meanwhile, I'm also seeing greater convergence in the AI debate at the practical level, particularly when it comes to the standardization of technology. These developments don't address every possible concern, but they are steps in the right direction.

When considering our regulatory approach, it's useful to draw parallels with the rise of the automobile in the early 20th century. After all, AI's current position is similar to that of the car in the 1920s: a technology that's been proved to work, but which still presents legitimate safety concerns. Well, throughout the 20th century, regulators responded to these dangers with bumpers, seatbelts, stop lights, and other traffic regulations. Even today, cars are not entirely safe – but a constant process of regulatory refinement has led to huge reductions in the number of traffic accidents, and the overall utility of the technology is no longer in doubt. I think the same applies to AI: we have a powerful but largely unregulated technology hitting the market, and over the coming years, we'll have to ensure it serves the public good.

Learning as you go is an important part of any technological revolution, and always has been. That's why I'm not overly concerned by what's going on in AI today – in fact, I'm excited. We're currently on a long, fascinating path toward embedding this technology on a wide scale. History shows us these paths are never straightforward – they often require organizations to completely rethink their processes – but history also shows us that real innovation is possible. If we remain focused on regulating the narrow AI technologies that are already transforming our societies, and avoid sensationalist narratives about impending AGI, I think the future of artificial intelligence is very bright indeed.

A portrait of a man with dark hair and glasses, wearing a blue suit jacket over a white shirt. He is looking directly at the camera with a neutral expression. The background is dark and textured, with a diagonal light-colored shape on the left side.

“We need to
look more
critically at
how – and **why** –
we use **technology**
in **our societies**”

Technology and the Future

Ilyaz Nasrullah

Ilyaz Nasrullah is a digital strategist with a background in software engineering. At his consultancy firm, Rumified, he specializes in the strategy, management, and implementation of innovation within established corporations.

In the early 2000s, there was great optimism about the benefits that digital tools and social media would deliver, with many people expecting they would advance democracy and social progress around the world. A few decades later, however, we're increasingly aware of the ways in which digital technology seems to erode democratic values and trust in society. Instead of assuming that technological innovation and social progress are one and the same, we need to look more critically at how – and why – we use technology in our societies.

This is where artists, writers, and other creatives can make a vital contribution. Take science fiction: the more media we consume about technological dystopias in which machines end up controlling our lives, the more likely this dystopian



future becomes! On the other hand, stories about social, not technological, advancement offer new inspiration; stories like *Dune*, for instance, which features a civilization that rejected artificial intelligence in the distant past. We should never rule out technological development, but first we need to understand the social vision our technology is meant to support. A new type of science fiction could certainly help us imagine a different technological future.

To be clear: the computer is an amazing innovation. It helps people bring ideas into reality. As someone with a computer science background, I understand why people are fascinated with them. But – because we associate technology with progress and because it's easy to look for tech fixes instead of addressing society's real problems – we've come to rely on tech more than I believe we should. Instead of maximizing efficiency and comfort with a tech push, we need to ask what a successful world looks like and what qualitative impacts we need to create to get there, and only then look at how technology can support that movement.

I don't think a tech-centric outlook is inevitable. While it's true that maximizing efficiency and throughput has been ingrained in the field of computer science from the beginning, the technology itself isn't the only reason these concepts are so influential. The pull of the "economic man" (*Homo economicus*) model – which suggests people are rational agents always looking to maximize their utility – has also played a part. Western and Western-influenced societies are still very driven by this view of human motivation and behavior, and it continues to influence a great deal of technological research.

At this year's ACM Conference on Fairness, Accountability, and Transparency (ACM FAccT), the winning paper was titled "The Values Encoded in Machine Learning Research". It focused on the motivations given by academics for doing their research in the first place. The finding was that relatively few papers justify how their project connects

to a societal need, and that only 1% discuss the negative potential of the project. Instead, values like performance, efficiency, generalization, and quantitative evidence were much more frequently cited as justification for the research being undertaken.

It's worth asking whether values like these truly advance innovation – or, indeed, stifle it. By moving quickly, we create unintended consequences. Much of the time, therefore, people are using their talents to fix the problems caused by technological developments, rather than to imagine and realize a better future. As a result, we're missing a direction for our technological development. The EU's declaration on digital rights and principles goes some way to filling this gap by outlining a value system and regulatory path for the use of technology in society. Even if it's not yet clear how many of its proposals will be implemented, I think it's a step in the right direction.

Consider new technologies like virtual reality (VR), augmented reality (AR), and mixed reality (MR), all of which are set to become mainstream in the next few years. Benelux tech organizations are well positioned to succeed

as these industries develop: the region has a highly trained workforce, a good digital infrastructure, and strong creative industries all going in its favor. But any involvement in immersive technology should focus on the six themes of the EU's declaration: people at the center, solidarity and inclusion, freedom of choice, participation, safety and security, and sustainability. For companies like Ordina, a unique value proposition could be to take a more measured approach to immersive technology by focusing on its qualitative and long-term effects.

This measured approach is essential for improving our relationship with digital, which is why I'm encouraged that new voices are now being heard in the debate around technology and society. As different perspectives gain greater representation in the media and the industry, there's growing awareness of the need to give people with different backgrounds and fresh ideas the chance to work on their vision, instead of just trying to fix some of the flaws in the current paradigm.

For me, this is an exciting shift, because we've reached a critical threshold when it comes to assuming that social progress and technological progress are the same thing. If we continue to make this assumption, we run the risk of creating unforeseen problems and doing damage to our societies. Now is the time to look beyond efficiency, quantification, and surveillance. If we can do that, and if we can put society before technology, a more fulfilling digital future is possible.



It's easy to look for
tech fixes instead of
**addressing society's
real problems**

A close-up portrait of a middle-aged man with dark hair and glasses, smiling slightly. He is wearing a dark jacket over a light blue and white striped shirt. The background is dark and textured. On the left side, there is a dark brown diagonal shape containing orange text.

“As a perennial
optimist,
I think we’re only
just beginning to
understand the **full
potential of what
technology can do**”

Technology and the Future

Peter Hinssen

A serial entrepreneur, Founder of nexxworks, keynote speaker, and author, Peter Hinssen is one of the most sought-after thought leaders on radical innovation, leadership, and the impact of all things digital on society and business. He lectures at various business schools, including the London Business School (UK) and MIT (Boston, MA), and his most recent book, *The Phoenix and the Unicorn*, explores how traditional companies can thrive in the digital age.

We're currently living in what I call the Never Normal: a world where seismic shocks are constantly triggering new developments. This ever-present change dynamic brings challenges that aren't just technological, but also biological, social, geopolitical, ecological ... the list goes on. Nevertheless, I believe technology can play a powerful role in tackling every one of these Never Normal challenges. In fact,




as a perennial optimist, I think we're only just beginning to understand the full potential of what technology can do.

Discussions of technological innovation often betray dystopian undertones, with people worrying that robots are going to steal our jobs and making other dire predictions. But if we look at an industry like healthcare, we see quite a different story. Historically, the sector has been one of the slowest to adopt digitalization – but today, all kinds of technological innovation and especially data science, are completely transforming healthcare, making it more affordable, more efficient, and more accessible.

I think we'll be the last generation to know so little about our bodies. Smart devices allow us to collect more and more data that can be used to improve patient outcomes. At the moment, medicine is very much a doctor-centric world, but technology can turn it into a truly patient-centric world, in which data enables the complete personalization of care. Our generation can leverage the power of innovation in healthcare – a shift that has the potential to make everyone's lives longer, healthier, and more enjoyable.

Developments like these show that technological innovation is no longer the preserve of the tech sector. We're now entering a phase where technology is making its presence felt in every aspect of every business, which presents huge opportunities for those willing to seize them. This applies not only to the Unicorns that have emerged in recent decades, but also to traditional companies that use the power of technology to become stronger and more innovative. I like to call these companies Phoenixes, and I think the coming years hold a lot of promise for them.

Walmart is a great example. It had been written off by many as an old-school retailer, bound to be eaten alive by Amazon. But its transformation has been truly spectacular. The company is reinventing itself as the retailer of the future, blending in-store experiences with e-commerce to become



Organizations need to ask how they can leverage the speed and power of a network to deliver results

more relevant for its customers.

On the other hand, when digital transformation is done poorly, companies simply take their old analog processes and put a digital veneer on top, hoping it'll scale up. The façade is bound to crack eventually, because there's no underlying digital transformation, just a digital translation. Real digital transformation is about fundamentally understanding how the organism of an organization needs to adapt – including its strategy, structure, culture, and processes. That means recognizing the business issue at heart, as well as the systems that need to be fundamentally altered to make digital solutions work.

It follows that a Phoenix company's ability to rise from the ashes comes down to culture. In many businesses, the cultural vector is sidelined as an HR responsibility, but I think that's a short-sighted view. You're not going to win in the Never Normal just by equipping your IT departments with a huge budget and hiring smart data scientists. You need to focus on the cultural dimension – on fostering the

qualities that will enable your organization to thrive in an ever-changing world.

What are these qualities? For me, velocity, agility, creativity, innovation, networking, and experimentation are key. Together, I call them the VACINE framework. The last two – networking and experimentation – are worth exploring here in more detail.

By networking, I mean recognizing that our world is no longer about siloes, but about reaching out. In this age of networks, leading companies are always looking to form partnerships, platforms, and ecosystems. Organizations need to ask how they can leverage the speed and power of a network to deliver results.

Experimentation, meanwhile, means taking risks – even if we fail. I've always loved this quote from Nelson Mandela: "I never lose. Either I win or I learn." Companies have plenty to gain from such a mindset. After all, in the Never Normal, there's no yellow brick road to success; you have to create your own path. To foster a climate of innovation, we need to build psychological safety into organizations in a way that empowers people to feel comfortable with taking risks. Start-ups naturally have an easier time of this: often, taking big risks is the only way they can succeed. But for traditional companies looking to become a Phoenix, nurturing that thirst for experimentation is equally vital for future success.

In many ways, the concept of the Phoenix applies to Ordina, which is on its own cultural transformation journey from IT supplier to digital business partner. The company already offers an essential service, but, more excitingly, it's challenging itself to experiment – and thereby become even more relevant to its customers in the Never Normal landscape. After all, providing hardware, software, or skills is one thing, but having a real impact on a customer's business is quite another.




“In addition to reducing energy consumption, organizations need to build **circularity, modularity, and repairability** into digital hardware”

Technology and the Future

Jacqueline Cramer

Jacqueline Cramer is Professor of Sustainable Innovation at Utrecht University, an orchestrator of circular economy transition initiatives in sectors including construction and textiles, and has been actively involved in sustainability issues since the mid-1970s. She was Minister of Housing, Spatial Planning, and the Environment in the Dutch government from 2007 to 2010.

Since 1970, human beings have been consuming resources faster than the Earth can replenish them. As things stand, we need 1.75 Earths to support our levels of demand – and by 2050, if we don't change our behavior, we'll need three. The problem is we only have one planet. To safeguard our world for future generations, we need to achieve carbon neutrality and build a circular economy. Digital technologies, if used responsibly, can play a vital role in helping us reach these goals.



For too long, organizations have taken a siloed approach to sustainability, and this has limited our ability to understand our overall impact on the environment. The good news, though, is that digital technologies can break down many of the barriers that stand in the way of large-scale collaboration. We can now share data in ways that were never possible before, empowering organizations throughout the value chain to hold each other accountable and become much more prudent in their use of resources. Through the Internet of Things (IoT), for example, we can use smart sensors to gather precise pollution data to in turn inform our sustainability efforts. Meanwhile, visualization technologies, such as virtual reality (VR) and augmented reality (AR), can help us reduce the environmental impact of construction projects by enabling planners to use advanced digital modeling in place of building materials.

It's easy to forget, however, that all these technologies consume vast amounts of energy. Today's consumers have a greater awareness of the energy they use at home, but tend to overlook the fact that their cloud computing is powered by energy-intensive data centers, often located halfway around the world (people are often shocked to hear that downloading a large video uses the same amount of energy as a return trip from Amsterdam to Rome by car!). Unless the digital sector does something to address this issue, it risks becoming the 21st-century equivalent of the coal industry.

Sharing server space is one way to reduce energy consumption at data centers. Companies tend to overestimate how much data storage is required for their operations – resulting in more servers than we need. LEAP, an initiative by the Amsterdam Economic Board, aims to tackle this by increasing the amount of shared data storage, and the city has also announced plans to fine data centers that don't impose "sleep mode" on idle servers in their facilities. Policies like these are an important step toward making cloud computing more sustainable.



Instead of encouraging
a race to the bottom,
governments need to
draw inspiration from
organizations that are
**constantly working to
raise the bar**

In addition to reducing energy consumption, organizations need to build circularity, modularity, and repairability into digital hardware. This is especially important when it comes to rare earth metals, which are used to manufacture essential components in computers, smartphones, and other digital technologies. The first issue is scarcity: unless we use our equipment for as long as possible and retrieve rare earth metals from products at the end of their life cycle, future generations won't have the same building blocks we use in digital technology today. The second issue is that rare earth metals are only found in certain parts of the world – meaning supply is closely tied to geopolitical developments. The war in Ukraine has already highlighted the dangers of resource dependency, and these problems will intensify

in the digital sector if we don't find a more sustainable approach.

All these problems are global in scale – and they require global solutions. Governments have a crucial part to play in these efforts, of course, but I'm also a big believer in the importance of front-runners: organizations that know what it takes to deliver sustainability improvements within digital technology. These frontrunners can show governments the key drivers for scaling up the whole innovation system, and governments can draw on their best practices to design effective policy instruments that complement what front-runners have already been doing. In my view, this approach is far superior to governments' current approach to sustainability, which often focuses on regulating companies that are lagging behind.

As important as individual front-runners are, no one organization holds the solution to making digital technologies more sustainable. What we need are "coalitions of the willing": organizations working together to improve the footprint of the IT sector and show their peers that bigger, bolder steps can be taken. In the years to come, Ordina can play a significant role in these coalitions by continuing to make progress on its ESG journey. The company is already a strong performer on the social aspect of ESG, especially with its various educational outreach programs.

If enough organizations unite around this common goal, it's possible for the IT sector to become carbon negative by 2040. With sustained efforts to reduce our energy consumption, increase our focus on circularity, and leverage the power of the latest digital tools, we can build an economy that respects the limits of the only planet we have. This future is within reach – but we need to work together to make it a reality.



“People **don’t complain** about digitalization if the new system is actually **more efficient!**”

Technology and the Future

Marcel Levi

Marcel Levi is President of the Executive Board at the Dutch Research Council (NWO), an organization that funds thousands of top scientific researchers. He is also Professor of Medicine at the University of Amsterdam and University College London (UCL).

All too often, the public conversation about healthcare is focused on problems, such as the lack of money, lack of professionals, long waiting lists, and other associated challenges. But when I look around my field today, I see a huge amount of positive innovation with the potential to change the conversation. After all, as a doctor, I always prefer to think in terms of solutions – and digital technology is certainly creating a solution revolution for the medical sector.

I'm extremely excited, for instance, about developments in artificial intelligence (AI) and machine learning. These techniques have a variety of promising applications in healthcare, including interpreting medical images, planning radiotherapy, looking at histopathology, or conducting



electrocardiograms. Using traditional methods, all these activities are labor intensive, but with the help of machines, they can be carried out much more quickly and with far greater precision. I don't believe AI will replace radiologists, radiotherapists, or pathologists, but this kind of tech will make medical work much more efficient.

I see the same trend when it comes to collecting and processing patient data. Data collection – on everything from delirium to bedsores – is a huge part of the healthcare profession today, and nurses can spend hours logging this information at computers located far away from their patients. At UCL, though, we found a simple digital solution to this problem: tablets. The new electronic health records work beautifully, allowing nurses to input data at the patient's bedside and spend more time interacting with people. What's more, it turns out that people don't complain about digitalization if the new system is actually more efficient!

The impact of digitalization can also be seen in telemedicine, which, after years of slow progress, has been hugely accelerated by the pandemic. Digital consultations save patients time on travel and in the waiting room, while doctors can help more patients per hour. I must stress, however, that telemedicine isn't the optimal solution in every case. When I have to break bad news to a patient, I always prefer to do that in person. But I now conduct between 30% and 50% of my consultations on digital platforms. As in many other sectors, digital tools are enabling hybrid solutions across the medical sphere – combining the best elements of both human and digital interaction.

Nevertheless, not every innovation in the medical world is a positive step towards addressing real problems. I think the value of certain solutions, like diagnosis support for doctors, is currently overhyped: as things stand, computers simply cannot make an effective diagnosis. It's interesting to speculate about why exactly this is; for instance, I don't think the human body is as schematic as many of these models

assume, and there are so many additional, qualitative factors to consider, such as nonverbal communication. In light of this complexity, I'm skeptical that computers will be up to the job any time soon: there's simply too much information that the models don't consider.

There's a similar flaw in some healthcare apps. Many are released with the promise of helping patients manage their asthma, heart disease, diabetes, or other chronic conditions. But on closer inspection, these apps often turn out to be solutions looking for a problem. The software is designed by very smart people, of course, but the providers of these tools are often unable to resolve the patient's problem, or address the issue with the healthcare system. Robotic surgery is another example: it's increasingly being introduced without sufficient evaluation and used in situations where it's not required. This unthinking embrace of technology is exactly what we need to avoid. Instead, we must keep our focus squarely on patient outcomes and follow whichever method will improve them the most.

Indeed, some of the biggest recent improvements in the medical sector aren't technological in nature. Human-focused solutions, such as giving healthcare professionals greater autonomy, can make an equally large impact. I've always found it strange that, after spending around a million euros training a surgeon, we don't trust them to make decisions. But I'm hopeful that this attitude is changing. At UCL, nurses, doctors, and other healthcare professionals displayed great initiative throughout the Covid-19 outbreak, for example by moving respiratory support normally reserved for intensive care units to the general wards. The idea came not from upper management, but from nurses and support staff. If management can recognize when to play a supporting role, rather than a leadership role, frontline health workers can use their experience and expertise to deliver better care.



We have to remain realistic, and not allow ourselves to get carried away on a **wave of tech optimism**

Why do I mention this example in the context of technology? Because it's important to remember what we're trying to achieve. As medical professionals, we want to do the right thing for our patients. Any solution, technological or otherwise, has to be assessed by that standard. If not, we'll lose sight of what made digital useful in the first place, and continue to design flashy apps that don't address real problems. The means by which the solution comes – whether it's digital transformation, cultural change, or any other route – is far less important than the solution itself.

That's not to say I don't believe technology will revolutionize the healthcare sector and, ultimately, transform it for the better. On the contrary, we'll see solutions emerge in the future that are now considered impossible. But we have to remain realistic, and not allow ourselves to get carried away on a wave of technological optimism. Some solutions will take longer to implement than others. Others won't be useful at all. But if we take all these things into account, and always ask ourselves the simple question of "What problem am I solving?", I think the future of medical technology looks very healthy indeed.