International Center for Future Generations ICFG

Five Emerging Technologies to Act On Now

International Center for Future Generations

Jakob Graabak, Maria Koomen and Max Reddel

Contents

Executive Summary

From the steam engine to electricity to the internet, technology has had a profound impact on the modern world. In the coming decade, five powerful emerging technologies are poised to fundamentally reshape societies.

These technologies—advanced artificial intelligence (AAI), neurotechnology, biotechnology, climate intervention technologies, and quantum computing and communication— have the potential to impact the lives of billions of people around the world, for better or worse. These five technology areas are also experiencing periods of large-scale investment prompting breakthroughs in capabilities and accessibilities. Further, they could rapidly develop in mutually reinforcing ways, presenting a particularly complex set of challenges and opportunities for society.

If harnessed for the common good, these technologies can help tackle some of the greatest challenges of our time, including climate change, disease, poverty, and more. However, without public oversight, they could lead to rampant inequality, political instability, economic upheaval, disregard for civil liberties, and threats to human life and ecosystems across the planet.

Whether these innovations are used more for good or more for harm depends to a large extent on the degree of public oversight over their development. Currently there is very little, because just a small number of primarily commercial actors control the bulk of the intellectual property, research, and development in these fields.

To limit these technologies' harmful effects, they must be governed by democratic principles, responsibly and with public accountability. The European Union (EU) is well positioned to lead in this field. It is an international actor with the protection of human rights and democratic values — as well as support for economic development — at its very core, and has recently passed several pieces of landmark technology legislation. All this puts the spotlight on the next EU political cycle in 2024-2029, which will present an opportunity to ensure new technology is deployed while respecting the public interest, public safety, democracy, the rule of law, and human rights.

This paper reviews the current state of the aforementioned five technologies and the potential risks they pose in the absence of responsible governance. Each section then offers pointers on policy questions and ideas for EU members and institutions to take into account in their strategic planning. Here is a brief overview.

Neurotechnology

Neurotechnology—devices that can read from and write into the brain—could address currently incurable medical conditions, from restoring lost eyesight to enabling paralyzed patients to walk again. It does this by using the most intimate human data, the contents of our minds. As neurotech becomes more refined and widely available, it could also be used to boost productivity or athletic performance. But technologies that steer perception, behavior and thinking open up avenues for manipulation or surveillance of the most dystopian kind imaginable. The EU must quickly address these new ethical and legal challenges by closing any loopholes in existing laws and how they're enforced, and by adopting new laws where needed.

Biotechnology

New applications of biotechnology, particularly gene-editing, genome-writing and computational biology tools, are used in everything from drug discovery to agriculture innovation. Rapid cost declines in recent years have supercharged their evolution. However, mistakes or deliberate misuse could release new pathogens into the wild, triggering pandemics much deadlier than COVID-19. Legislation must tighten control of DNA supply chains, laboratory safety, and research projects. Furthermore, the EU and international partners must strengthen disease detection, transmission prevention, and rapid-response capabilities across the world.

Climate Interventions

With the world on track to exceed 1.5 degrees Celsius of global warming, researchers are exploring tools to artificially cool the planet very quickly. Climate intervention technologies such as solar radiation modification (SRM), also known as geoengineering, may buy the human race crucial time to decarbonize the global economy and adapt to inevitable climate impacts. However, many uncertainties remain and the potential risks could harm ecosystems and the ozone layer, and trigger global conflict. The EU should take a proactive role in supporting, shaping and promoting the international governance of climate interventions.

Quantum Computing and Communication

Quantum computers in particular promise to help solve problems that today's computers cannot, in fields ranging from energy to logistics. But when sufficiently powerful quantum computers become commercially available, hackers could use them to strike at the security of the financial system and other critical infrastructure, leading to chaos on a global scale. The EU must develop comprehensive technical standards for future-proofing information infrastructure against attack.

Advanced Artificial Intelligence

AI has already transformed modern life and will touch almost every industry in the years to come. But governance of these systems lags far behind the pace of their evolution. Even AI engineers cannot explain their systems' behavior. In this context, the EU's AI Act is a groundbreaking first step towards establishing binding oversight and accountability for the developers of leading AI systems. However, it may not be sufficient to curb the risks of AI systems being developed in the near future, given the rapid pace of technological development in this field. With no current scientific consensus around how to develop AI systems safely, the EU needs to invest in a science of safe AI, and to plan for contingencies in case current efforts fall short.

This new wave of emerging technologies obliges the EU to move even faster and more effectively to govern technological progress within the rule of law and with an eye towards human well-being. Given the rapid pace and unpredictable consequences of these innovations, and the global geopolitical and regulatory landscape, the EU must take a proactive approach to secure the safety, human rights, and democratic principles of citizens now. While the five technologies unpacked here are considered the most urgently in need of governance action, there are other areas that the International Center for Future Generations is monitoring and evaluating for future efforts.

Recent measures like the AI Act are important but just the beginning. Implementing, enforcing, updating, and future-proofing such regulations will be a never-ending task. And even as it performs this task for AI, the EU needs to ensure robust frameworks along similar lines for domains such as neurotechnology and climate intervention technologies. The EU and like-minded partners can pass laws and regulations that encourage research, development, and innovation while still protecting public safety and individual agency. Understanding the risks associated with these technologies and how they intersect with one another is an essential first step.

Introduction

Five powerful technologies emerging in the coming decade promise to both enhance and disrupt society. **Artificial intelligence** (AI) systems can teach themselves new skills without explicit training, giving them unpredictable capabilities. AI-driven **neurotechnologies** can read from and write into the brain directly, a boon for some but raising the specter of surveillance or mind control. **Biotechnologies** that can engineer life and design new drugs—or new diseases—are becoming both more powerful and more accessible. **Climate interventions**, such as solar radiation management, could quickly cool the atmosphere but might destabilize ecosystems. **Quantum technologies** could tackle once-intractable computational problems and create super-secure communications, but also break current forms of encryption.

As the world's economic and technological powers compete to out-innovate one another, better and more agile governance is needed.

These technologies are developing fast—in some cases exponentially. They can be used in many different ways, which makes their effects hard to predict. And they will mutually reinforce one another, making the effects even more unpredictable.

What also sets these technologies apart is that a small number of actors, often in the private sector, are developing and controlling them, without adequate public input or accountability.

This creates an uncomfortable knowledge and talent gap between the technologies' developers and public institutions. The gap makes it hard for regulators to react fast to new developments. Given the profound implications for human agency, identity, and well-being, civil society and academia need to step up and help regulators get ready.

This paper zooms in on five technologies that have the potential to go mainstream within the next decade. They are already often advancing faster than scientists and engineers expected. And AI, aside from being one of these five technologies, will also turbocharge the other four, making them both more powerful and more accessible. The benefit to societies and ecosystems could be enormous. The technologies' unfettered development and use, however, could fuel inequalities and destabilize democratic societies and economies.

As the world's economic and technological powers compete to out-innovate one another, better and more agile governance is needed to ensure these powerful technologies do more good than harm. Future-proof regulatory frameworks take years to design, adopt, and enact. Without holding back innovation, public institutions must constantly update their understanding of new developments and their implications.

The European Union (EU), as an international, rights-driven actor, is in a position of special responsibility to lead societies through the coming era, especially given its unmatched track record for legislative and regulatory action that addresses complex global issues. European Commission President Ursula von der Leyen has appealed to European leaders to lead the way on a new global framework for AI, based on the EU's AI Act, the world's first comprehensive and horizontal AI regulation. But with many more technological advances on the horizon, a holistic approach to technology governance is needed.

A new EU leadership term starts in 2024. The International Center for Future Generations urges political parties and decision-makers to make the governance of emerging technologies a priority in this new term.

Understanding the Emerging Technologies

Here is a summary of the five powerful emerging technologies and their implications that EU institutions ought to take into account as they plan the next term. While the five technologies unpacked here are considered the most urgently in need of public oversight, there are other areas that ICFG is monitoring and evaluating for future efforts.

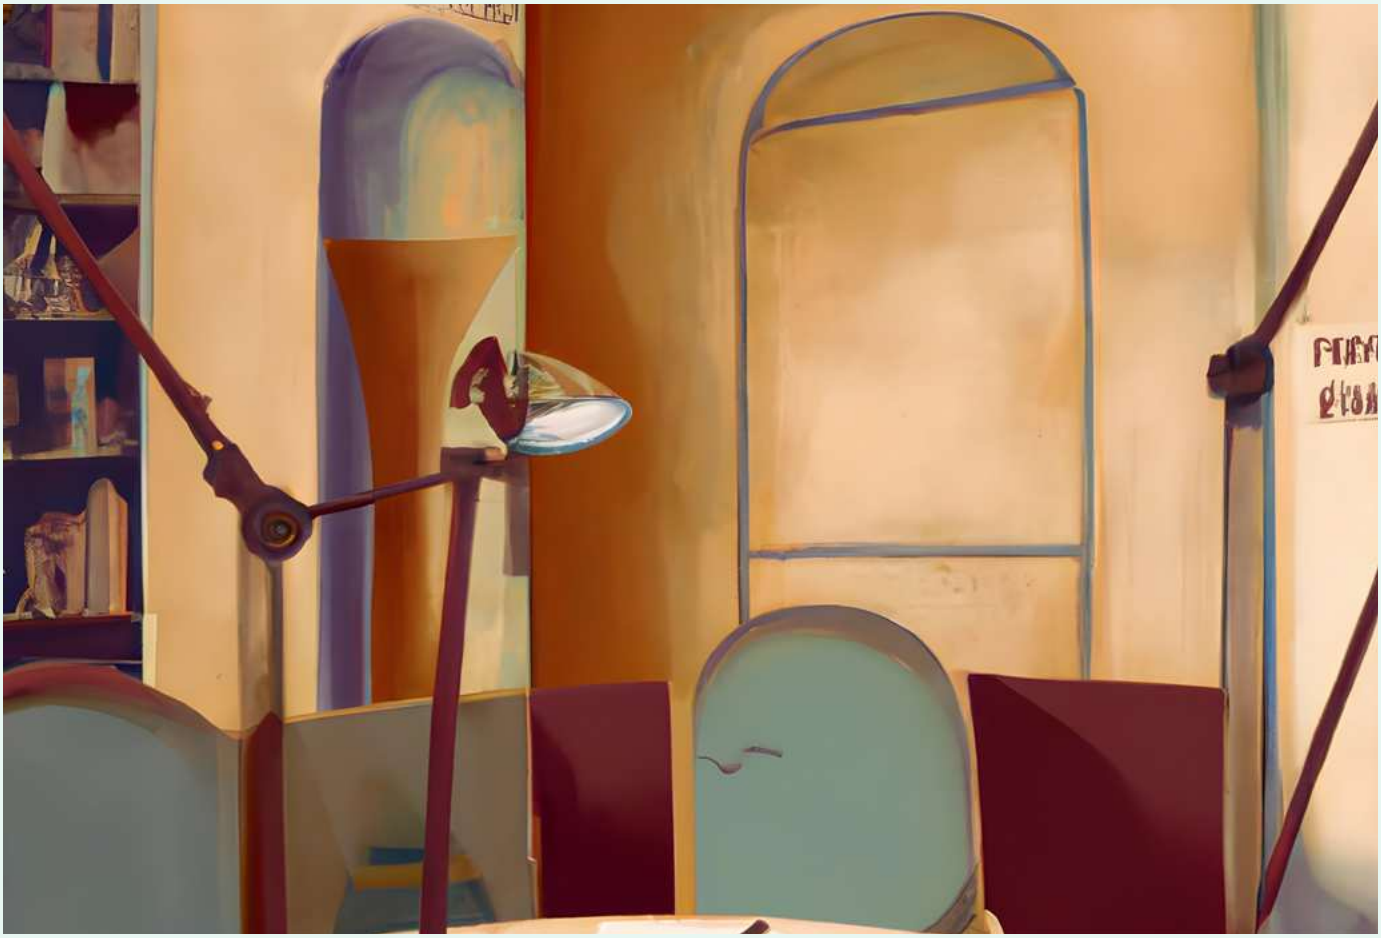

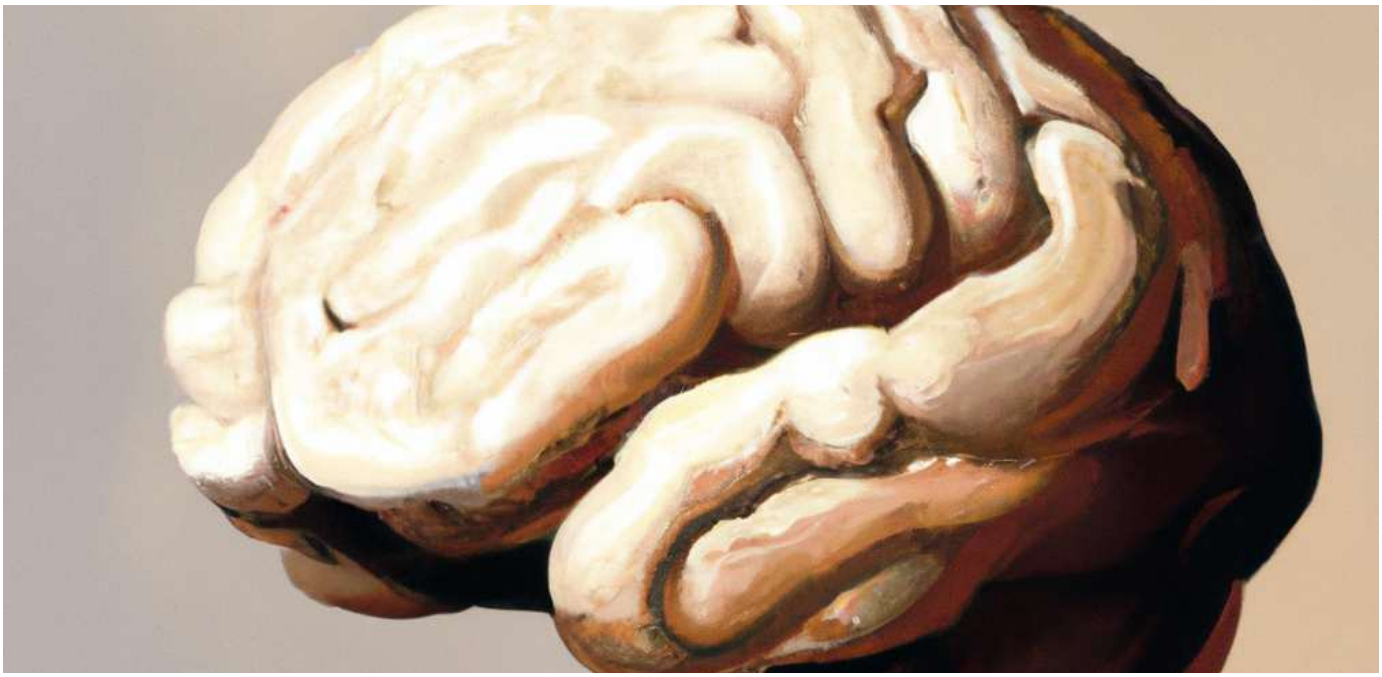

01 Neurotechnology

The State of Play

Neurotechnologies, or technologies that can read from and write into the brain, are the next frontier of human enhancement. Scientists anticipate that in five to 10 years, medical applications of these technologies could restore sight to blind people, help those who have been paralyzed to walk again, and enable patients with severe paralysis to communicate. As quickly if not sooner, wearable or implantable consumer devices could enhance learning abilities, boost physical performance for sport or combat, treat insomnia, reduce anxiety, and supercharge productivity. The consumer market may eventually drive developments in brain-controlled devices. But this also means people's brains—and the cognitive data that can expose their most private thoughts, feelings, preferences, habits, health, conscience, decision-making, and reasoning—could be up for grabs.

Potential Implications

In the consumer market, devices could be hacked or contain so-called neuro-cookies—similar to the cookies that track web users' browsing habits—to identify the wearers' consumer or political preferences, or even implant new ones. This would give advertisers new ways to target and influence consumers, as well as giving political actors more power to sway voters through indoctrination or interference. Meanwhile, unevenly distributed access to these devices—inevitable with any new technology—could further increase social and economic inequality.

Unconstrained use of neurotechnology could also infringe civil liberties. Law-enforcement agencies could subpoena data from the devices to examine citizens' thoughts and use them as evidence of a crime—or even of the intent to commit one, as in the film Minority Report. This would extend the use of mass surveillance to the cognitive realm and spark uncertainty about the role of intent in judicial systems. If a state were to criminalize certain beliefs or intentions, then "thought crimes," as imagined by George Orwell in 1984, could become provable and prosecutable, undermining the most fundamental right to freedom of thought. Meanwhile, militaries are already developing neurotech devices to support decision-making and targeting processes in combat, which may raise questions

People's brains and the cognitive data that can expose their most private thoughts, feelings, and preferences can be up for grabs.

as to whether human fighters are fully in control and accountable.

These technologies encroach on the intimacy of the human brain, hitherto taken for granted. Scholars and practitioners from a range of disciplines have flagged the risks neurotechnology poses to the rights to privacy, freedom of thought, freedom of expression, and self-determination. Further, as neurotechnology converges with AI, the harms and risks of AI may be transferred to neurotechnologies, making them prone to algorithmic biases, race and gender disparities, and big data threats, such as cybersecurity and profiling applications.

All this raises ethical and political questions that should be settled before these technologies become widespread. Does the possibility of thought control outgrow existing definitions of fundamental rights? Are so-called neuro-rights needed? Do existing data protection laws adequately safeguard the brain? What are the moral and ethical limits of consumer applications in neurotechnology? Are open societies and democratic systems prepared for the proliferation of mind-reading and thought-controlling technologies? And are open markets suited to address dual-use concerns and the risks of biases and inequities?

Governance Options

Experts and regulators all over the world are starting to address these implications. There has been exploratory dialogue at the international level, and standard-setting efforts at the multinational level. Chile has proposed new legislation and Argentina, France, and Spain are working on their own versions. These efforts are important groundwork, but more is needed.

The EU, building on its legacy of respecting fundamental rights, including in the digital domain, should seek a leading role. And given how soon neurotechnology could go mainstream and how profound the impacts could be, there is no time to lose.

To start, the EU could engage in impact scoping and scenario planning on ethical, legal, and cultural questions. How powerful can these technologies become? Beyond their intended uses, what could their benefits and harms be? What are the limits to their use for mitigating human disabilities or augmenting human capabilities? Where should red lines be drawn on future development and use? And are decision-makers ready to denounce unacceptable research, development, or use?

Then, EU member states should collaborate to identify gaps in existing laws and rights frameworks, and shape guidelines for product standards. The EU will need to decide between three approaches: preserve the existing situation, where neurotechnology devices are treated like any other medical or consumer device; make special provisions for neurotechnology within existing laws and regulations such as the AI Act, the Medical Device Regulation, and the General Data Protection Regulation (GDPR); or issue neurotechnology-specific regulation to reflect the area's unique and complex nature.

Whichever of the three approaches the EU takes, member states and institutions must also make sure existing laws are enforced. Policymakers should consider new mechanisms of oversight and accountability to avoid loopholes that could leave human rights and democratic values vulnerable. National and EU authorities must have sufficient budget, expertise, capacity, mandate, and—not least—political will for enforcement, while patients, consumers, and civilians alike must have sufficient access, avenues, and agency to file complaints and lawsuits.

02 Biotechnology

The State of Play

Biotechnology is the scene of some of today's most profound technological advances. New techniques make it possible to measure and manipulate life with unprecedented precision at every level, from individual molecules up to cells, whole organisms and entire ecosystems.

Falling costs have played a big part here. For instance, genome sequencing became more than 180,000 times (or 99.999 percent) cheaper from 2001 to 2022. By comparison, solar cells, known for their steep cost declines, became only 23 times cheaper over the same period.

The applications are legion. Accurate protein-folding prediction using algorithms like AlphaFold allows scientists to quickly zero in on potential new drugs. Precision gene-editing with tools like CRISPR-Cas9 can be used to eliminate congenital diseases or make crops more resistant to parasites. Gene drives, which spread new genes through a population as its members reproduce, could be used to alter or even eliminate an entire species, such as mosquitoes. Biologists are learning to write new genetic code from scratch, allowing them to—for instance—create yeasts that naturally synthesize pharmaceutical drugs.

Other applications include mRNA vaccines, which played a crucial role in accelerating the end of the COVID-19 pandemic; novel antibiotics to combat antimicrobial resistance; cell and gene therapies with the potential to cure chronic diseases; lab-grown meat and meat replacements; and new treatments for cancers and other non-infectious diseases like obesity. In total, experts estimated in 2020 that biotechnology could unlock $2–4 trillion of annual economic value in 2030–2040, linked to better human health and performance, agriculture and food systems, and consumer products and technologies.

Potential Implications

Currently, developing dangerous pathogens still requires access to biological raw materials, specialized lab equipment, and advanced expert knowledge. However, all three barriers are gradually eroding or already compromised.

Genetic materials can be ordered online from a range of DNA providers. Do-it-yourself biotechnologies, such as next-generation benchtop DNA synthesizers, are lowering the need for expensive lab equipment to manipulate DNA with precision.

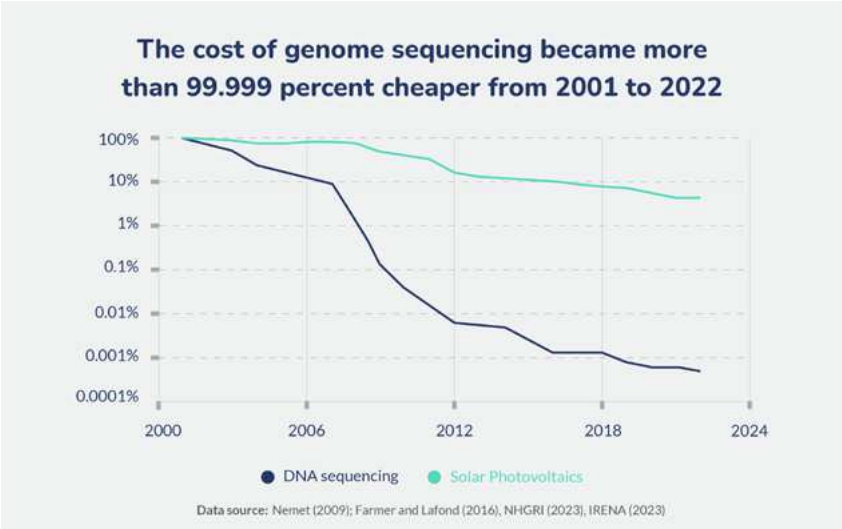

Wetterstrand KA. DNA Sequencing Costs: Data from the NHGRI Genome Sequencing Program (GSP) Available at: www.genome.gov/sequencingcostsdata. Accessed 20/02/24

International Renewable Energy Agency (2023); Nemet (2009); Farmer and Lafond (2016) – with major processing by Our World in Data. Available at: https://ourworldindata.org/grapher/solar-pv-prices. Accessed 20/02/24.

Detailed instructions for building viruses have been assembled for nearly all viral families, reducing the need for domain expertise. These instructions are meant to accelerate medical research but could also serve as blueprints for malicious or reckless actors.

If this trend continues, which is likely—especially with the help of advanced AI—engineered pathogens could become the main source of pandemics in the future, either deliberately created by malicious actors or accidentally released by researchers or engineers working without proper safeguards.

This kind of risk is relatively new. While terrorists have attempted—with occasional success—to use bioweapons in the past, technical difficulties have mostly deterred them. These difficulties barely hold today and certainly won't in the future.[1] As for accidental releases, there have been many such incidents in high-risk laboratory work over the last five decades, and the number of biolabs working with contagious pathogens at the highest safety levels has boomed in recent years.

Decisive action is needed today to ensure that pathogens capable of starting new, deadly pandemics cannot be created and released.

Yet despite these threats, there is neither a global early-warning system against emerging pathogens, nor the medical supplies and response capabilities needed to contain an outbreak once detected, as COVID-19 demonstrated.

Future pandemics could be much worse than COVID-19. The deadliest outbreaks in history, including the Black Death in 14th-century Europe and smallpox in the Americas, may have killed more than 50 percent of the population at the time, though estimates vary. By comparison, the estimated excess mortality from COVID-19 was around one in 300 people globally, and even the 1918–1920 flu pandemic and the two world wars are estimated to have killed single-digit percentages of the global population.

An engineered pathogen of similar destructive potential as smallpox or the plague might not kill as many people, thanks to modern medicine, but COVID-19 showed that a disease with even a relatively low mortality rate can be hugely disruptive to a complex, interconnected society. Decisive action is needed today to ensure that pathogens capable of starting new, deadly pandemics cannot be created and released.

Governance Options

COVID-19 showed that the world cannot currently detect and contain pandemic threats before they spread globally. When the tools and knowledge for engineering pathogens are much more widely available, a global detection-and-response infrastructure will be needed. It must be able to contain pathogens more transmissible than even the Omicron strain of SARS-CoV-2 and to detect and defend against new biological threats, even if competent malicious actors are trying to circumvent the defenses.

Until these global detection and defense mechanisms are in place, stronger governance of DNA supply chains, laboratory safety and security, and dual-use research could help mitigate the risks from engineered pathogens. Today, DNA providers

[1] Most famously, the Japanese cult Aum Shinrikyo, with the help of geneticist and early member Seichii Endo, used sarin gas on the Tokyo Metro in 1995 and attempted to weaponize Ebola a few years earlier, while the Rajneesh cult succeeded in contaminating several salad bars in the United States with salmonella.

> “Current trends suggest that within a decade, tens of thousands of skilled individuals will be able to access the information required for them to single-handedly cause new pandemics.”
>
> **Kevin Esvelt**
> MIT Biologist

representing roughly 80 percent of the market are members of the International Gene Synthesis Consortium, committed to voluntarily screening their orders to limit harmful use. Screening needs to expand to cover the other 20 percent.

This could entail extending screening to new geographies, embedding screening mechanisms in multinational collaboration—through the World Health Organization, for example—and further strengthening screening practices. One way to do that would be to require DNA synthesis providers to take out liability insurance, providing a market incentive to identify and eliminate screening loopholes. Screening should also be extended to next-generation benchtop DNA synthesizers to catch smaller-scale malicious actors.

Advanced AI could also lower the bar to engineering novel pathogens.[2] Leading AI executives have testified in the US Congress that the next generation of these tools, which may be available in just two to three years, could offer substantial assistance to bad actors.

However, these actors may not even need help from AI. An area of scientific research called pandemic virus identification (PVI) aims to assess which pathogens can cause pandemics, in the hope of accelerating progress on vaccines and other treatments. Such knowledge would also serve as a blueprint for a weapon of mass destruction. In his 2022 testimony to the US Senate, Massachusetts Institute of Technology biologist Kevin Esvelt stated that “successful [PVI] will immediately cause widespread proliferation” and that the risks to public health of such discoveries would outweigh the benefits by a factor of 100. For this reason, PVI and other concerning areas of dual-use research need to be strictly governed. This should include global restrictions on funding and publishing research that might lead to widely accessible bioweapons.

Still, these kinds of restrictions on supply chains, lab safety, and dual-use research will not prevent engineered pandemics in the long run. They can only buy time while a global detection system is set up. This system will need to be able to identify a broad range of emerging pathogens, not just the most likely culprits, and it will have to cover all major international airports and other sources of transnational transmission, as well as equipping medical centers with the resources to detect shifting disease patterns.

Global defenses must be able to contain a new disease outbreak within its first few days, through a combination of pharmaceutical and non-pharmaceutical measures (like ring vaccination around outbreaks and targeted quarantines). Also needed is better infrastructure to block transmission, such as indoor ventilation, air filters, and disinfecting lights. Finally, essential workers will need transmission-free supply chains for critical goods like food and medical supplies, so society can keep running if containment fails and the new pandemic spreads out of control.

All this requires significantly more investment in pandemic mitigation than in the past. Countries with developing healthcare systems will especially need help improving their detection and rapid-response capabilities. In the wake of COVID-19, the EU has already started future-proofing its pandemic mitigation through the European Commission’s Health Emergency Preparedness and Response Authority (HERA). Two years in, HERA is investing in a broad range of new technologies that can help protect Europe from emerging biological threats.

Now is the time for EU policymakers to commit the resources and formulate a bold, targeted biodefense agenda. Beyond saving untold lives and livelihoods in the event of a new pandemic, these investments would bolster public health in general. Reducing incidence of the seasonal flu, common cold, and new strains of COVID-19 could alone save thousands of lives and billions of euros a year.

[2] A common counterargument is that some of this information is already available in modern search engines. In the future, more capable large language models and biological design tools could plausibly aid malicious actors in ways that traditional search engines cannot. Attempts to install guardrails against this kind of misuse have failed so far.

03 Climate Interventions

The State of Play

The year 2023 was a watershed moment for climate change. It was the warmest year on record by a wide margin. Yet, despite intensifying extreme-weather events, the world is not decarbonizing in line with the 2015 Paris Agreement. Global average temperatures are set to rise far beyond the goal of 1.5 degrees Celsius above preindustrial levels. Even if all the national pledges from Paris are fulfilled—a big if—a rise of 2.5–2.9 degrees by century's end is likely. Many scientists warn tipping points are imminent, marking irreversible change in the Arctic, the Amazon, and elsewhere.

On the current warming trajectory and under current policies, some countries will see double-digit percentage-point increases in mortality and potentially multi-decade economic recessions[3]. The World Bank estimates that some 216 million people may migrate by 2050.

In addition to climate tipping points, experts warn of social tipping points: moments when populations affected by acute climate disasters put pressure on policymakers to take drastic action. These are especially likely in the most vulnerable countries, which may experience lethal heat waves as soon as 2030.

Such drastic action could include solar radiation modification (SRM), also known as geoengineering. This involves increasing the reflectivity of the atmosphere to bounce more sunlight back into space so that less heat reaches the planet's surface. Several methods have been proposed, of which stratospheric aerosol injection (SAI) is by far the most researched and discussed. It would mean injecting aerosols into the stratosphere which would then spread globally, reflecting sunlight in much the same way as the dust from a large volcanic explosion.

SAI would require specialized aircraft. Existing planes might be modified to carry out a stunt deployment of SAI, but not in a scientifically robust way on a global scale. One expert interviewed estimated that it would take more than 10 years before the planes and infrastructure needed are operable.

[3] However, estimates of the economic impacts of climate change differ greatly, from 0–5 percent impacts at 3 degrees Celsius global warming, according to some papers, to 15–30 percent, according to others.

> “Given that SRM would affect all countries if ever used, the lack of international oversight seems negligent at best, and enormously dangerous at worst.”

Cynthia Scharf
Practice Lead,
Climate Interventions,
ICFG

Potential Implications

If deployed at scale, SRM and particularly SAI would affect every country in the world, but not equally. The uneven distribution of impacts may prompt disagreement about the optimal global temperature, and some countries may perceive (correctly or not) that SRM is worsening their local climates and seek redress. These claims would be hard to adjudicate, since current science makes it impossible to demonstrate beyond doubt which effects were attributable to SRM and which would have occurred anyway.

Moreover, once SRM has begun it needs to be sustained for decades. If it were abruptly halted after several years, temperatures would rapidly shoot back up, which would be disastrous for many ecosystems. SRM should only be phased out after emissions have reached net zero and massive amounts of atmospheric greenhouse gases have been removed. Any SRM program would thus need a remarkably stable, long-term international governance regime.

Even a controlled, agreed deployment carries risks. The Intergovernmental Panel on Climate Change indicates that these risks include potential harm to the environment and to social, political, and economic systems as well as to geopolitical stability. There are also ethical, moral, legal, and intergenerational equity and justice issues.

On the other hand, not deploying SRM means a much higher risk of catastrophic climate change as the world shoots past the 1.5 Celsius benchmark and heads toward double that figure. And as long as the world delays adopting SRM, the chances grow that domestic political pressure in one or more climate-vulnerable countries will push them to unilaterally launch a stunt deployment. This would probably not be enough to significantly cool the planet but would nonetheless send geopolitical shockwaves and potentially trigger conflict. In short, there are no risk-free options.

Governance Options

SRM is currently not governed by any comprehensive international framework, and multilateral climate diplomacy takes years to develop. It took 30 years of negotiations before governments said they would transition away from fossil fuels at last year’s UN climate conference, COP28. SRM might be operable at scale in a decade; the time to start developing international guardrails for it is now.

All governments must have access to the latest science, and be a part of governance discussions, but it’s especially crucial that the most climate-vulnerable countries have a strong voice. They have the most to either lose or gain (we don’t even know enough yet to say which) if SRM is ever deployed at global scale.

As Cynthia Scharf, practice lead for climate interventions at the International Center for Future Generations, argues, “As climate impacts intensify, there is a growing risk that countries or even non-state actors might rush to use SRM. Given that it would affect all countries if ever used, the lack of international oversight seems negligent at best, and enormously dangerous at worst.”

The EU has already established itself as a forward-looking, responsible global actor on climate change. It has led by example to help meet the Paris Agreement temperature targets and has adopted ambitious climate laws to get there. In June 2023, the European Commission called for global talks on climate intervention technologies, saying their risks, impacts and unintended consequences are poorly understood and “necessary rules, procedures and institutions have not been developed.”

The EU now has the opportunity to put this call into action—both within the EU itself, in developing guidelines for research and indoor testing, and through multilateral institutions such as the UN Environment Program, the World Meteorological Organisation, the UN Framework Convention on Climate Change, and the Montreal Protocol. Governance for climate interventions needs to be set within the context of climate policy overall and not as a stand-alone. Events such as the UN’s Summit of the Future in September 2024, or other intergovernmental arenas, should also consider including conversations on climate intervention governance.

Whether one opposes, supports, or is unsure about climate intervention technologies, there is no question they need to be governed before someone uses them. The riskiest path is ignorance and inaction.

04 Quantum Technologies

The State of Play

Governments and institutions worldwide, including the EU, are investing heavily in quantum research and development. A total of $55 billion of government investment has been recorded worldwide. The EU has invested €1 billion ($1.1 billion) in its Quantum Technologies Flagship, a long-term research and innovation initiative that aims to "put Europe at the forefront of the second quantum revolution." In December 2023 several member states signed a quantum declaration with the aim of making Europe the "quantum valley" of the world and the "leading region globally for quantum excellence and innovation."

Quantum technologies harness the laws of quantum physics to offer unprecedented capabilities. They encompass three main categories: sensing and metrology, communication, and computing and simulation.

Quantum sensors can measure motion, space, and time more precisely than legacy devices and could transform a broad range of fields, from energy and transportation to healthcare, finance, and security. Quantum communication technologies encrypt data or form the basis of a so-called quantum internet that could connect quantum computers globally. Quantum computers may be able to solve currently intractable global problems beyond the capacity of today's computers, for example in materials science, mathematics, energy, and logistics. Quantum computing research is already pushing the limits of information processing and developing new machine-learning methods, improving research and development, and optimizing supply chains.

These uniquely powerful technologies could accelerate progress in crucial areas, such as the UN Sustainable Development Goals, according to experts at the new

[3] However, estimates of the economic impacts of climate change differ greatly, from 0–5 percent impacts at 3 degrees Celsius global warming, according to some papers, to 15–30 percent, according to others.

Once quantum computers are powerful and reliable enough, they are expected to be able to crack the encryption that underlies practically all forms of cybersecurity and data protection today.

Open Quantum Institute, "by enabling the development of more efficient drugs, cheaper fertilizers, longer-lasting batteries and more efficient solar panels."

Quantum computers have been in development for decades, but recent progress has been exponential. For example, in December 2023, IBM launched the first 1,000-qubit quantum chip—qubits are the equivalent of transistors in an ordinary digital computer—after years of releasing chips that doubled the number of qubits every year. These computers, though, are still unstable and notoriously prone to error. Error-correction techniques are being developed that compensate for this by, essentially, treating hundreds of qubits as one and averaging out their results. However, this means the machines need to be much larger to do useful work. Scientists estimate that a computer would need to have millions of physical qubits to deliver the equivalent of thousands of error-corrected ones.

Potential Implications

Once quantum computers are powerful and reliable enough, they are expected to be able to crack the encryption that underlies practically all forms of cybersecurity and data protection today. Some industry experts expect this to happen as soon as 10 to 15 years from now. These computers "could crack the cryptography that underpins financial stability," according to the International Monetary Fund, rendering banks, payment systems, Europe's digital single market, currency values, and the global financial system vulnerable. Journalists, human-rights defenders, and for that matter anyone who communicates or stores information online, will also be at risk of attack or unlawful surveillance.

It's time for the EU itself to start forging a proactive, coordinated quantum cybersecurity strategy.

"Quantum-safe" or "post-quantum" encryption standards are already being developed. Nonetheless, if critical infrastructures don't transition to these new security standards in time, there could be large-scale hacking and information leakage. Anticipating this, bad actors are reportedly already gathering encrypted information in so-called harvest now, decrypt later attacks, for example on banks and currencies.

Still, given the economic opportunities, there will likely be a commercial and geopolitical race for quantum dominance similar to those today in microchips and AI. This means a new global race for information security and safety as well.

Governance Options

The transition to quantum-safe information security systems requires several elements: quantum-safe standards, algorithms, crypto-agility (the ability to switch the cryptographic algorithm used), new technologies (for example, quantum key distribution, which relies on a physical distribution method rather than an algorithmic one), new and faster update systems, more knowledge and expertise, and enormous coordination—especially for regional players like the EU and currencies like the euro.

While the EU is contributing to research and innovation in quantum technology, it hasn't yet tried to shape technical standards, despite the European Commission's vow to take a "more muscular approach" to standardization. EU leaders must ramp up the legal, technical, and political capacity to face and guide the quantum revolution, domestically and internationally.

A handful of EU member states, including Denmark, France, Germany, and the Netherlands, are developing or implementing national strategies for post-quantum encryption. It's time for the EU itself to start forging a proactive, coordinated quantum cybersecurity strategy. Importantly, for the EU to lead the world in setting standards, any proposals that would weaken encryption should be off the table.

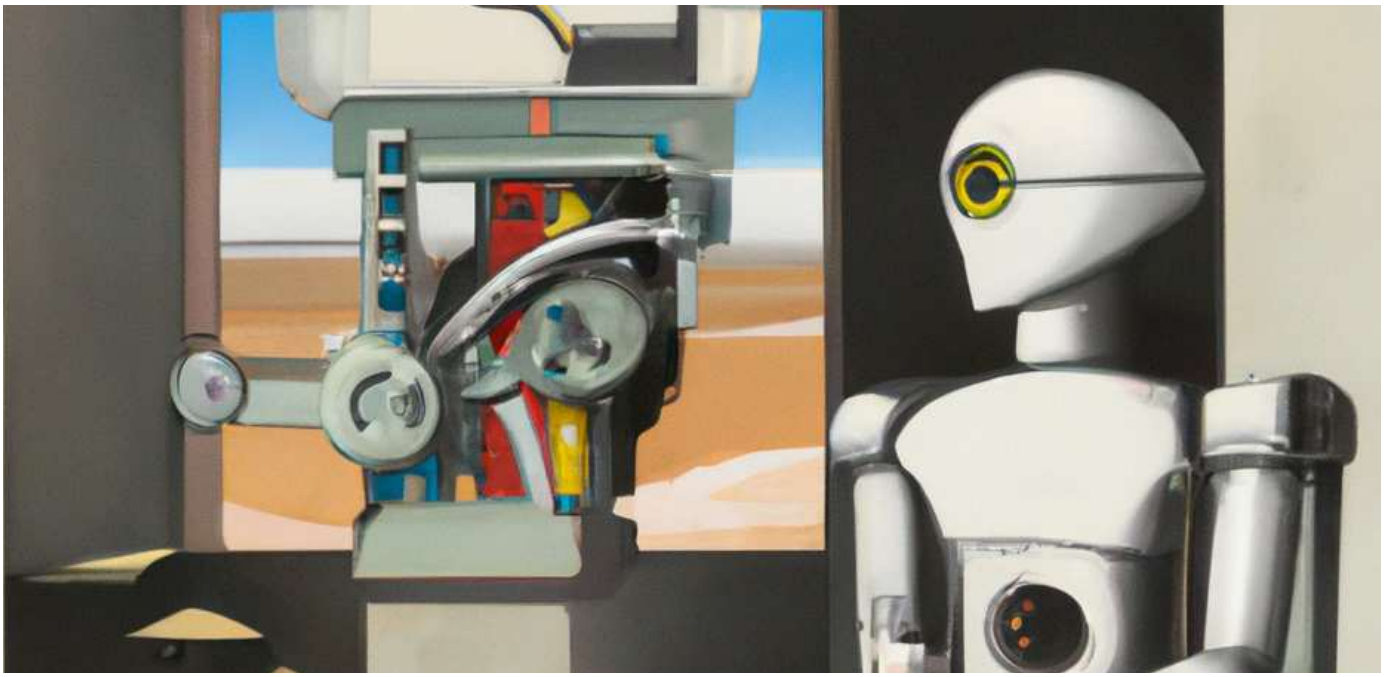

05 Advanced Artificial Intelligence

The State of Play

Over the past decade, AI has transformed from a promising technological frontier into an all-purpose technology that is reshaping industries, societies, and lives. In the coming years, it is expected to boost progress in many fields, including brain medicine, biology, climate sciences, and physics—and to accelerate progress in the four other emerging technologies mentioned above.

At the same time, experts, activists, and policymakers alike are raising a broad range of concerns linked to AI. Possible dangers include the amplification of biases and inequality, automated misinformation, addictive products, unlawful surveillance, bioweapon design, cybercrime, and large-scale accidents.

The commercial engineering of increasingly capable systems greatly exceeds any scientific understanding of those same systems. This in turn hinders a well-informed public debate, fact-based public policies, and the evolution of governance models that ensure AI is accountable to the public.

(A definitional note: "Advanced AI" refers to the most capable general-purpose systems—foundation models trained on broad data, capable of a wide range of downstream applications, and powerful enough to pose severe risks to the public interest, global security, democracy, and fundamental rights. Advanced AI stands in contrast to narrowly designed AI systems that can do specific tasks such as create playlists of songs that match the tastes of individual users or determine mortgage rates that reflect the profiles of individual applicants.)

Modern advanced AI has matched or surpassed human-like competence in a broad range of standardized cognitive tests, including language comprehension and image recognition. Soon, it will become proficient and increasingly reliable in domains like mathematics and code generation. These

models are trained to excel at one specific task: anticipate the next word in a sentence by recognizing language patterns. But some models have shown emergent capabilities they were not explicitly trained for, like strategic planning or understanding aspects of the physical world. At the same time, even the highest-performing models are variable and unpredictable, and confidently make mistakes and "hallucinate" (i.e., make stuff up).

The progress in advanced AI since 2010 has been driven largely by an increase in the computational power (or "compute") and data used to train AI models. The compute needed to train leading models in 2010 corresponded to around one minute of thinking with a human brain, while modern models use the equivalent of more than 2,000 years of brain time.[4] However, scaling has failed to eliminate problems with safety, fairness, and reliability. Leading AI companies admit that they still don't know how to address these issues even as they build and push to market increasingly advanced AI systems.

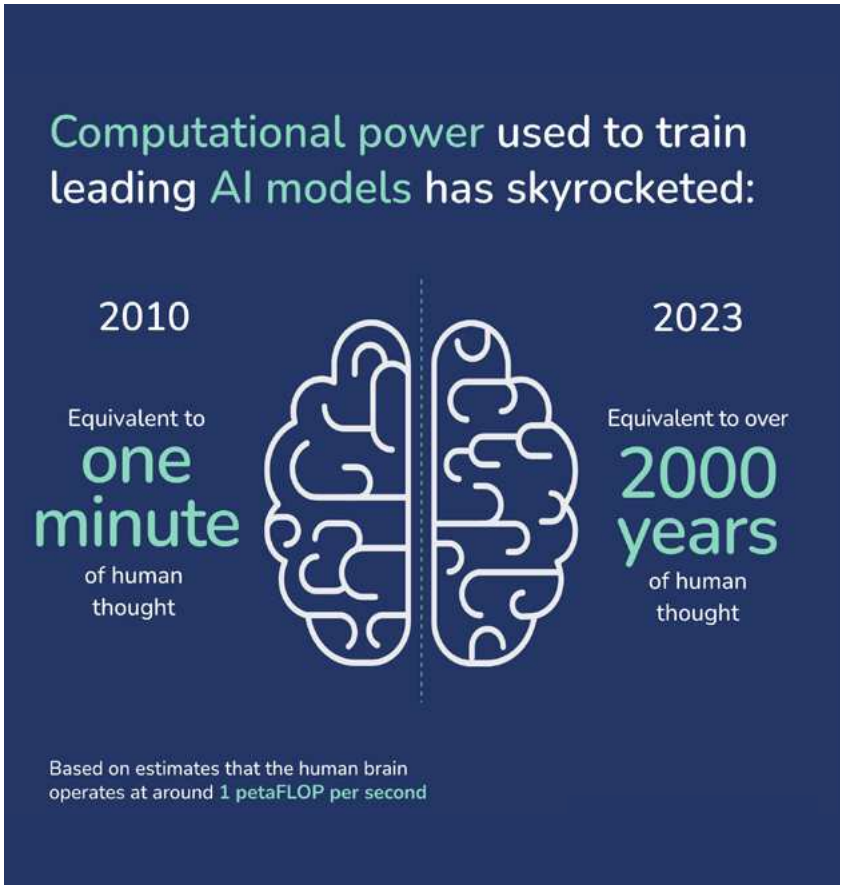

Data source: Epoch AI via Our World in Data

Instead of being coded step by step, these systems are grown and shaped by countless iterations of feedback and patching until they mature and perform—hopefully—as desired. This makes it very hard for developers to predict their machines' behaviors or guarantee any limits to what they will do: they simply don't understand the systems well enough to make such predictions with confidence.

Potential Implications

Experts disagree on three core questions about the future of advanced AI: which new capabilities will emerge; what their societal implications will be; and how hard it will be to mitigate the biggest risks.

Although recent progress in AI has been driven by growth in data and compute, continued scaling may not yield returns at the same rate. Instead, capabilities may plateau. Alternatively, they could progress more slowly but steadily, contributing to continued economic growth—although the gains could be unevenly shared, especially if workers are displaced by machines faster than new jobs can be created.

On the other hand, genuinely transformative types of advanced AI could emerge, such as the much-anticipated artificial general intelligence (AGI), capable of performing nearly all tasks that humans can do at least as well. AGI could drive explosive economic growth and societal impacts. Geoffrey Hinton, one of the three founders of deep learning, believes that AGI may emerge within the next five to 20 years, whereas others argue that it will take longer. However, expectations are changing fast. A survey of nearly 3,000 AI researchers, published early this year, found that the median estimated timeline to reach human-level performance had dropped by 13 years since a similar survey conducted just 14 months earlier.

When it comes to the societal implications, some experts focus on the fact that AI can

[4] Based on estimates that the human brain operates at around 1 petaFLOP per second. This is a highly simplified and approximate estimate, as computation in human brains and digital computers works very differently. However, it can shed some light on how fast the systems have scaled in recent years.

Some have flagged concerns that increasingly powerful, hard-to-audit, and technically unstable models could break out of their makers' control.

exacerbate existing issues. It often perpetuates and amplifies societal biases contained in the data it is trained on. Tools based on such AI, such as hiring or criminal sentencing algorithms, have the potential to further marginalize minority groups. Additionally, AI may be harnessed for widespread unlawful surveillance by powering face recognition, voice transcription, and social-media analysis software. Authoritarian governments can employ these to monitor, predict, and influence citizen behaviors on an unprecedented scale. The AI-based algorithms underpinning social media platforms help amplify and spread misinformation, and generative AI tools such as chatGPT allow malicious actors to easily create fake text and images or simply mediocre, cheap content to power clickbait websites, all of which degrades the information ecosystem. The profits these systems generate have led to the monopolization of power and control among the handful of companies that lead their development.

Besides these direct harms, advanced AI systems may supercharge other sectors, such as synthetic biology. In the future, a malicious actor may be able to use advanced AI to design a novel pathogen, potentially killing billions of people. Advanced AI could also greatly accelerate scientific progress across fields such as quantum computing and autonomous weapons, potentially fueling new great-power conflicts.

Some have flagged concerns that increasingly powerful, hard-to-audit, and technically unstable models could break out of their makers' control.[5] In the direst scenarios, experts have warned that increasingly capable algorithms could take on critical decision-making without meaningful human oversight. The world could then be at the mercy of whatever these models are instructed to achieve and the unpredictable strategies they devise to do so—even if this leads to large-scale harm.

Governance Options

Last year saw a surge in governments' focus on AI risks and the development of policy responses at the G7, the AI Safety Summit, the OECD, and the EU. The EU's AI Act, as the first multilaterally binding initiative, is a pioneering step toward governing AI with a rights-based, risk-prevention framework. The Bletchley Declaration, from the November 2023 AI Safety Summit, highlights the need for international cooperation on safety risks associated with frontier AI, particularly highly capable foundation models. Complementing these supranational efforts, countries such as the UK, the US, Japan, Singapore, Brazil, and China are also developing national AI governance initiatives.

But experts disagree about how to harness the potential of these technologies while minimizing harm. One core uncertainty centers on what it would take to mitigate the potential existential harms from advanced AI, with proposals ranging from modest risk management and oversight measures to enforced pauses or even a permanent halt to the development of advanced AI systems. Governments sometimes fear investments will go elsewhere if policies are adopted, but this common view should be rejected. In fact, regulatory clarity on AI will enable companies to make use of new tools and could even open up new market segments.

Given the rapid developments in this field, decision-makers need to look to the future, and skate where the puck is going. This requires action on three fronts at once:

1. Establish reliable measures for oversight and accountability, such as monitoring and auditing frontier models and ensuring accessible redress in cases of harm.

2. Carry out scientific research on the societal implications of advanced AI to be able to make informed governance decisions.

3. Create contingency plans against proliferation in case it proves harder than expected to build ever more advanced AI systems safely.

[5] See for instance these open letters (1, 2), where the signatories include many of the leading academics who championed the breakthroughs underpinning current AI technologies, such as deep neural networks and the transformer architecture

Decision-makers need to create contingency plans against proliferation in case it proves harder than expected to build ever more advanced AI systems safely.

In terms of oversight and accountability, the EU's AI Act is a great first step that covers the entire AI product life cycle. It includes vital measures like risk assessment, model evaluations, and cybersecurity. Still, the act describes these measures only at a very high level. Their efficacy will depend on how they are implemented and how the act is enforced.

While the act mainly covers narrow AI applications, it also briefly mentions general-purpose AI, which technically includes advanced AI. This is a very important aspect of the act, since advanced AI systems will likely have the largest societal impacts and are therefore the most important to govern. It is essential that the EU AI Office embrace its mandate to keep the act relevant as advanced AI systems become increasingly powerful. For this to succeed, regulators will need to meaningfully oversee the developments of leading AI companies and anticipate their potential implications for European citizens and institutions, using the flexibility in the AI Act and developing new legislation as necessary.

As for the state of research into safe and ethical AI by design, nascent initiatives include national safety institutes in the US, the UK, and Japan as well as a State of the Science report commissioned by the UK. While laudable, these initiatives are not proportionate to the challenge. Based on current understanding, advanced AI could contribute to everything from utopian to dystopian futures, but these research efforts to safeguard the future count only a few hundred employees. The EU, which finances some €350 billion of research and development each year, could greatly increase the scientific understanding of how to build AI safely.

The efficacy of the EU's AI Act will depend on how its measures are implemented and enforced.

If advanced AI doesn't plateau but continues to get more powerful, policymakers will face the question of whether and when the harms will outweigh the benefits. Already, they are drawing red lines for specific domains where AI should not be applied—the "unacceptable" category of the AI Act. But this may not be enough to curb the harms of more general-purpose systems, especially if these become so advanced that they can deceive their users or resist attempts to control them.

It may then be necessary to put enforceable international anti-proliferation measures in place. This would require strong international collaboration and robust enforcement mechanisms. Few (if any) existing regulatory efforts are considering such contingencies, though 2023 saw a flurry of research on related topics. One idea is to use access to compute resources as a possible bottleneck to build enforcement around. Other research efforts have explored what international treaties or institutions for nonproliferation of advanced AI could look like.

It's too early to tell whether such strong action will ever be needed. But the discussion can and should begin today to outline which future scenarios would warrant a globally coordinated and enforced anti-proliferation regime for AI. This means that discussions of the future of AI cannot center only on the most likely scenarios but need to cover a range of possible futures and anticipate when and how to act in each case.

Conclusions

These five emerging technologies could shape the trajectory of the future. Their power and potential suggests that we are experiencing a pivotal moment in modern history. It's also, therefore, an opportunity for the EU to assert itself on the global stage by guiding technological progress toward a future that prioritizes stability, fundamental rights and values, and the well-being of humankind.

This requires a holistic approach to technology governance and will entail trade-offs. Rather than develop policy for each of the five technologies in a silo, the EU's approach should be integrated, recognizing that the technologies will influence one another.

The principles underlying existing technology policies—such as adaptable and tiered guardrails for safety, human rights, and other democratic principles, as well as responsible innovation—should extend to the newer domains.

The governance of these technologies will also need constant updating. That means investing in research and development, integrating foresight into political processes, fostering public dialogue, and developing flexible regulatory frameworks that can absorb rapid technological change. Regular reviews will be needed, both to seek updates on technological developments and to identify enforcement gaps.

Effective governance is not a curb on businesses but can instead give them a competitive edge.

Policymakers should also assume that technological advances will often outpace policy processes. Although EU lawmakers started work on the AI Act in 2021, the explosion in generative AI in 2022 and 2023 showed the limitations of the draft act before it was even agreed on. This was not a failure of policymaking, but the new normal.

Educating regulators, decision-makers, and the public about the potentials and risks of new technologies will be essential. Civil society, with the support of academia, has an active role to play in building a more robust science-policy interface to counterbalance the enormous influence of commercially driven actors.

Innovation can flourish under clear ethical and safety guidelines. Effective governance is not a curb on businesses but can instead give them a competitive edge. A predictable regulatory environment supports and attracts economic activity. It also boosts consumer confidence at a time of great political and economic uncertainty.

This paper does not prescribe over-regulation but recommends addressing areas of under-regulation, which can otherwise lead to distortions in the market and incalculable side-effects for society. Regulatory sandboxes are a useful tool for real-world testing of advanced technologies within well-defined guardrails. They allow for fine-tuning of legal boundaries in ways that empower small businesses and start-ups to innovate.

In addition to its domestic regulatory efforts, the EU must actively engage in international governance discussions and standard-setting. This could lead non-EU countries to adopt its governance approach, influence future technological advances on a global scale, and create a conducive environment for responsible innovators to conduct business worldwide.

In summary, the EU's role and interest in tech governance extends far beyond AI and far beyond the union's own geography. It encompasses a responsibility to steer the development of transformative technologies in the best interests of democratically governed societies. But the EU—or any other actor—can't do it alone. As the union ushers in new leadership, political agendas, and investment priorities over the next five years and beyond, it must prioritize global leadership and collaboration, anchored in the rights and values of future generations worldwide.

About the Authors

Jakob Graabak is the foresight lead for the International Center for Future Generations. He specializes in horizon scanning, technological progress, and scenario analysis.

Maria Koomen is the governance lead for the International Center for Future Generations. She specializes in emerging technologies, democracy, and governance.

Max Reddel is the program lead for advanced artificial intelligence at the International Center for Future Generations. He specializes in understanding technological advancements, societal impacts, and governance options for advanced artificial intelligence.

About ICFG

The International Center for Future Generations is a think tank dedicated to shaping a future where decision-makers anticipate and responsibly govern the societal impacts of rapid technological change, ensuring that emerging technologies are harnessed to serve the best interests of humanity.

Image credits

Generated with AI Shutterstock (reference to Giorgio de Chirico)

WWW.ICFG.EU