

COPENHAGEN
INSTITUTE
FOR FUTURES
STUDIES



INCLUDING 10 TRENDS
& ISSUES TO WATCH IN

2025

Exploring the futures of technology

UNDERSTANDING THE LAYERS IN AN AI-MEDIATED WORLD

By Copenhagen Institute for Futures Studies

This report explores the accelerating landscape of emerging technologies and provides a foundation for developing a futures-informed mindset. It shares insights into key trends and critical uncertainties of 2025, and look into the interconnected dynamics of technological advancements, with a focus on AI's emerging role as a mediator in the future. Its aim is to help readers make sense of the challenges and opportunities that lie ahead in 2025 and beyond.

The report is published by the Copenhagen Institute for Futures Studies (CIFS)

If you have any questions or comments, you are welcome to reach out to:

Futurist and Head of Media, Sofie Hvitved, shv@cifs.dk

Futurist and Head of Technology & Innovation, Bugge Holm Hansen, bhh@cifs.dk

Design & illustrations Sara Frostig/Firefly

© Copenhagen Institute for Futures Studies, January 2025

**COPENHAGEN
INSTITUTE
FOR FUTURES
STUDIES**



Content

How to read the report	4
Introduction	7
The need for organisational preparedness	8
The evolving geopolitics of tech	8
The AI-mediated framework	10
10 key trends & issues to watch in 2025	15
1: The complementary power of big data and small data	16
2: When context becomes king	18
3: Agentic AI – much more than chatbots	20
4: AI Clones & our own multiple representations	22
5: Bringing intelligence to the physical world with Physical AI	24
6: The rise of humanoid robotics	26
7: The rise of dual-use technology	28
8: The uncertainty of scaling	30
9: Environmental implications of an AI-mediated infrastructure	32
10: Quantum coming closer	34
Foresight in an AI-mediated world	36
Strategic future-driven questions to ask in 2025	39

How to read the report

How do you navigate a world where technological advancements accelerate faster than ever? Where AI plays an increasingly important role. What do you need to know and be aware of? How can you prepare for a world of constant change while remaining relevant in the future?

This report is designed to help you think critically about emerging technologies in 2025 and to consider their broader societal, economic, and organisational impacts in an increasingly AI-mediated world. Its purpose is not to offer definitive answers but rather to frame critical questions to consider as you unpack future issues of emerging technologies. It's about understanding how the different technologies evolve in an interconnected relationship that calls for new perspectives on how we understand and approach them. It's not about incremental, siloed advancements; it's about the connections that emerge when technologies collaborate across domains.

At the heart of the report is the **AI-mediated framework**. The model places the human at the core, surrounded by layers that represent AI mediated systems, their enabling technologies, the foundational infrastructure, and the broader worldviews that shape their context. This map is not just a visual framework – it's a tool to explore the relationships and opportunities that arise when technology is seen as part of a larger AI-mediated system.

The report then highlights **10 key trends and issues** to be watched in 2025. From the challenges of scaling AI and quantum computing to the transformative potential of contextual AI, agentic AI, and humanoid robotics, the selected trends and issues all bring significant opportunities while raising ethical and practical concerns as well. They underscore the complexities of balancing progress with responsibility.

For practical guidance, the report ends with considering how futures thinking and foresight itself can benefit from AI. One answer to the question of how organisations can work with the future in an AI-mediated world is that it requires equally **AI-mediated foresight toolbox**.

Finally, the **strategic questions** provide forward-looking prompts to inspire reflection and action, making them the perfect starting points for discussions within your organisational planning.

Read sequentially for context or skip to the sections most relevant to you. Use the critical questions throughout to provoke deeper thought and align strategies with future possibilities. This is not just a report meant to inform but a tool to guide proactive engagement with the technologies shaping our future.

Introduction

The pace of technological advances in recent years has left most of us caught between wonder and whiplash. On the one hand, we're witnessing incredible breakthroughs and hearing about even more on the horizon – advances that would have seemed like science fiction not long ago. On the other hand, it can feel like we're on a rollercoaster, struggling to catch our breath as each new innovation ripples through society faster than the last.

This isn't just about changes in technology; it's about the transformation of society at large. We're seeing industries redefine themselves from the ground up. AI has moved from research labs and sandboxes into our daily lives, while new technologies are starting to bridge the gap between human cognition and digital systems in ways that challenge our traditional notions of human capability and experience. The impact extends further than just faster computers or better algorithms – it's fundamentally changing how we live, think, approach problems, work creatively, and structure our businesses, organisations, and societies. Tomorrow's technological opportunities may seem endless, but grasping their true potential requires us to look past what we take for granted in the present and remain open to questioning our own understanding of the world while working inside an AI-mediated framework.

The opportunities brought by emerging technologies come with challenges and responsibilities as well. While businesses that are struggling to become future-ready may want to focus on radical, future-driven innovation, values such as openness, equality, fairness, and progress for the common good also need to be considered. How should we strike a balance between these different yet equally critical priorities?

The rapid advances in AI highlight this dilemma well. As AI is becoming ever more integrated into nearly every field of business and technology, its impact and influence are felt everywhere. Whether it's making content 'liquid' and allowing it to flow between formats and platforms, helping us rethink supply chains, delivering breakthroughs in health (like Google DeepMind's

AlphaFold Protein Structure Database that won the Nobel Prize in 2024), or reshaping education and the information ecosystem in general, AI tools for planning, research, diagnostics, and decision-making have never been more advanced and impactful.

The need for organisational preparedness

Technology has always been a tool for staying competitive and expanding into new markets. Even in traditionally strong business areas, AI will become an increasingly important tool for staying relevant in 2025. The risk of missing out on radical innovations that will transform industries in the future is significant.

This shift demands new organisational approaches. Many businesses remain rooted in systems designed for a different era, making it difficult to keep up with a shifting technological landscape. The challenge goes beyond simply adopting new tools; it requires rethinking how organisations operate, innovate, and adapt to an AI-mediated world.

This transformation comes with an increasingly steep learning curve and with a risk of turning chaotic if not managed properly. Skills gaps are widening, and there's growing pressure to prepare both the current workforce and future generations for an AI-driven world. This requires more than technical training; it demands building confidence and resilience to navigate uncertainty. In doing this, a balance must be struck between maintaining what works in the current system and making room for new ways of thinking and operating. Within organisations, there are often pockets of innovation where individuals or teams are already experimenting with how possible futures could look like. These efforts need to be recognised, supported, and connected. At the same time, it's important to acknowledge that change isn't easy. Many people and systems are deeply tied to established practices, and letting go of these can be difficult and uncomfortable.

It's not about quick fixes or waiting for technology to mature anymore. It's about having the courage to foster a culture of experimentation, adaptability, and learning – and about dedicating the resources needed to support it. Organisations that embrace uncertainty, welcome radical change, and commit to evolving alongside a shifting world will be the ones best equipped to thrive in the future.

At the same time, the push to advance AI capabilities while preserving human creativity, autonomy, and judgement introduces additional challenges. Balancing these ambitions with the practical demands of building competitive business models is no small task – and as a business or organisation, much depends on the specific market you operate in.

The evolving geopolitics of tech

The geopolitical race for tech sovereignty influences the direction of innovation and defines the environment in which businesses operate. While some regions, like the EU, prioritise ethical frameworks and data privacy, others push for dominance through aggressive investments in critical technologies. This creates a “battle of narratives,” where contrasting visions of the future

compete on the global stage. For example, the development of AI governance frameworks in the US, EU, and China is three distinct and competing narratives that reflect the broader ideological divides about individual rights and state control within these regions.

Narratives shape how we perceive the world, influencing beliefs, setting priorities, and defining realities. Narratives are powerful mental models used by governments, corporations, and institutions to shape public opinion and frame critical issues like climate change or AI regulation. Controlling the narrative has become essential for impact and shaping the future of how the technologies are being integrated.

The global competition in critical technologies – influenced by the competing narratives that drive divergent paths of development – underscores the need to anticipate disruptive scenarios. For instance, China’s deliberate long-term investment in quantum computing, artificial intelligence, hypersonic capabilities, and other areas of high-tech innovation has positioned it as a dominant force in 37 of 44 tracked critical technologies.¹ The EU is not dominant in any of them. While the EU forges ahead in its efforts to ensure user-centric regulation, China’s burgeoning technological supremacy risks undermining Western democracies’ influence in global standard-setting as well as their efforts to secure supply chains.

Simultaneously, the rising dominance of private entities in sectors like space technology presents both exciting opportunities and significant vulnerabilities – perspectives that shift depending on who is looking. From the private side, companies like SpaceX, with thousands of operational satellites in orbit, are driving innovation and offering transformative possibilities in communication, navigation, and security. However, this growing reliance on non-state actors also raises concerns about overdependence on private enterprises, potentially introducing risks that rival traditional geopolitical challenges.

Other crucial questions to consider are whether technological advancements will lead to a widening of global inequalities and whether it will help reinforce authoritarian control. There are also perspectives to consider that relate to the risk that unchecked resource competition will trigger geopolitical conflicts.

Questions such as these may be uncomfortable, but they demand honest reflection. As emerging technologies reshape society, their unchecked development risks magnifying the undesirable inequalities and issues we already face today. Ensuring future governance systems, in both the private and public sectors, that can address these power balances will be vital in order to shape our preferred future in an AI-mediated world.

The AI-mediated framework

In this report, we aim to shift the focus from the hype surrounding emerging technologies to the lens through which we can view their integration into a broader framework. In 2025 we see a move towards a world that is increasingly AI-mediated.

Being AI-mediated means that AI acts as an intermediary in a process, interaction, or decision. The AI analyses data, automates tasks, or provides recommendations, shaping outcomes or experiences. In the AI-mediated framework, the human takes centre stage, whether as a consumer, a co-worker, or an individual – making it more critical than ever to understand the individual needs, behaviours, and context.

The AI-mediated framework refers to a system where AI acts as an intermediary that facilitates, connects, or influences interactions, processes, and systems. In this context, AI is positioned as a key enabler, bridging gaps between different components – whether they are digital tools, physical environments, or human activities. This framework is made to highlight the role of AI in shaping relationships between individuals, technologies, and societal infrastructures, often focusing on its ability to personalise, optimise, and streamline these interactions. It is meant as a conceptual model to navigate how AI integrates into and impacts various layers of our technological and social ecosystems.

Imagine a world where you no longer need to navigate complex processes or systems to find what you need. Instead, AI provides personalised, seamless interactions tailored to an individual's preferences and needs, whether it's discovering products, accessing services and information at the workstation, or co-creating hyper-personalised experiences. These systems anticipate needs, adapt to the time and context of the user, and simplify decision-making. AI-mediated systems that simplify decision-making, adapt dynamically, and create a future where engagement feels effortless and deeply aligned with individual and collective priorities. AI will act as filters, cutting through the noise to deliver clarity and (personal) relevance.

In this reality, interacting without some form of AI mediation will become increasingly rare. As the AI-mediated ecosystems expand, not everyone will embrace them unconditionally. Counter-movements are already forming, driven by concerns over data privacy, the decline of critical thinking, and an over-reliance on algorithmic curation. Some worry this might create a kind of “mental obesity,” where we lose the ability to think deeply or independently. For these groups, opting out of AI-mediated interactions primarily represents a deliberate choice to preserve human agency and maintain deeper, unfiltered, and unmediated connections.

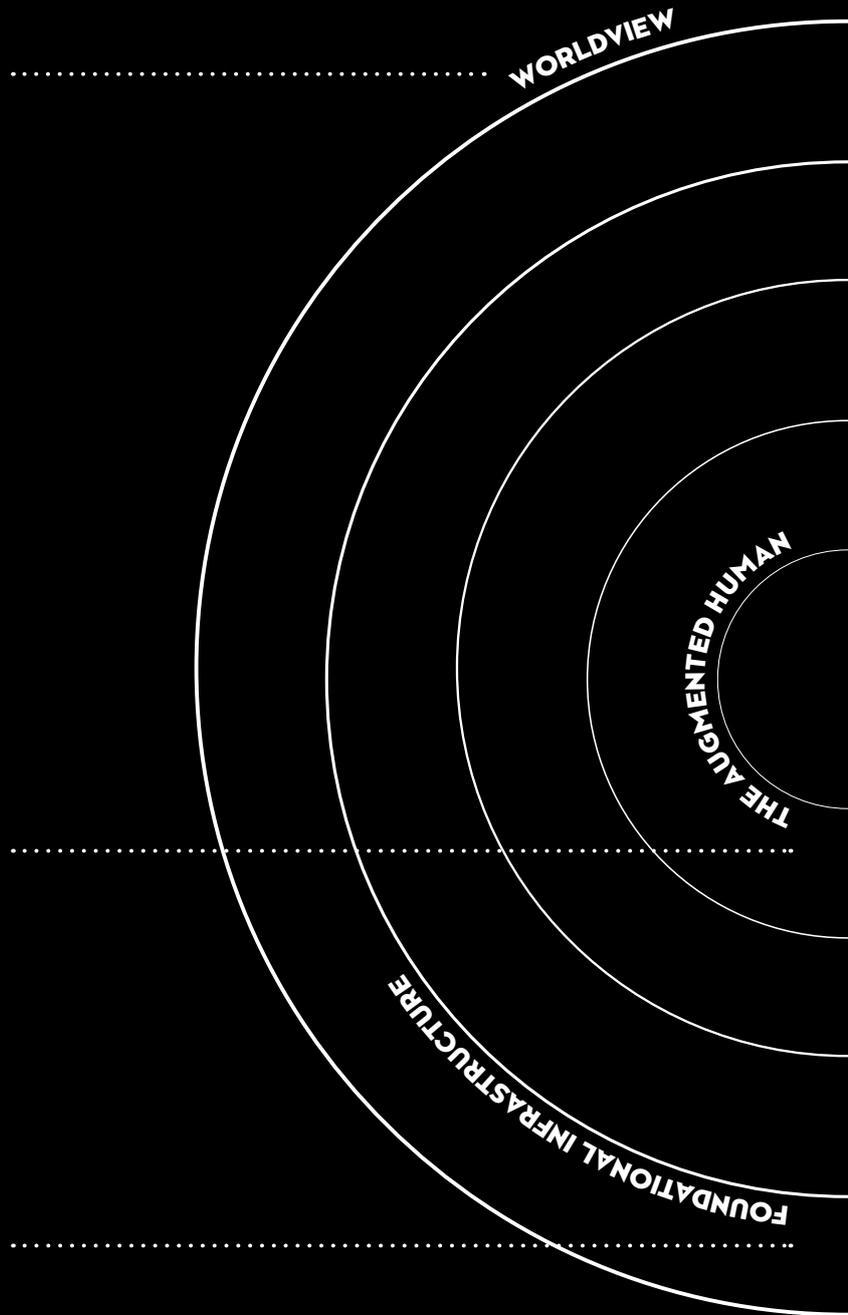
Despite these emerging counter-narratives, as seen historically, the majority are likely to prioritise convenience and personal relevance over moral considerations of the possible long-term impacts on humankind. The ease of hyper-personalised experiences, tailored recommendations, and seamless interactions often outweighs concerns about data sharing or the subtle trade-offs in autonomy.

This interconnectedness highlights the need to not view AI in isolation, but as part of a larger, evolving ecosystem mediating both the present and our possible futures.

“The AI-mediated framework refers to a system where AI acts as an intermediary that facilitates, connects, or influences interactions, processes, and systems.”

The AI-mediated framework

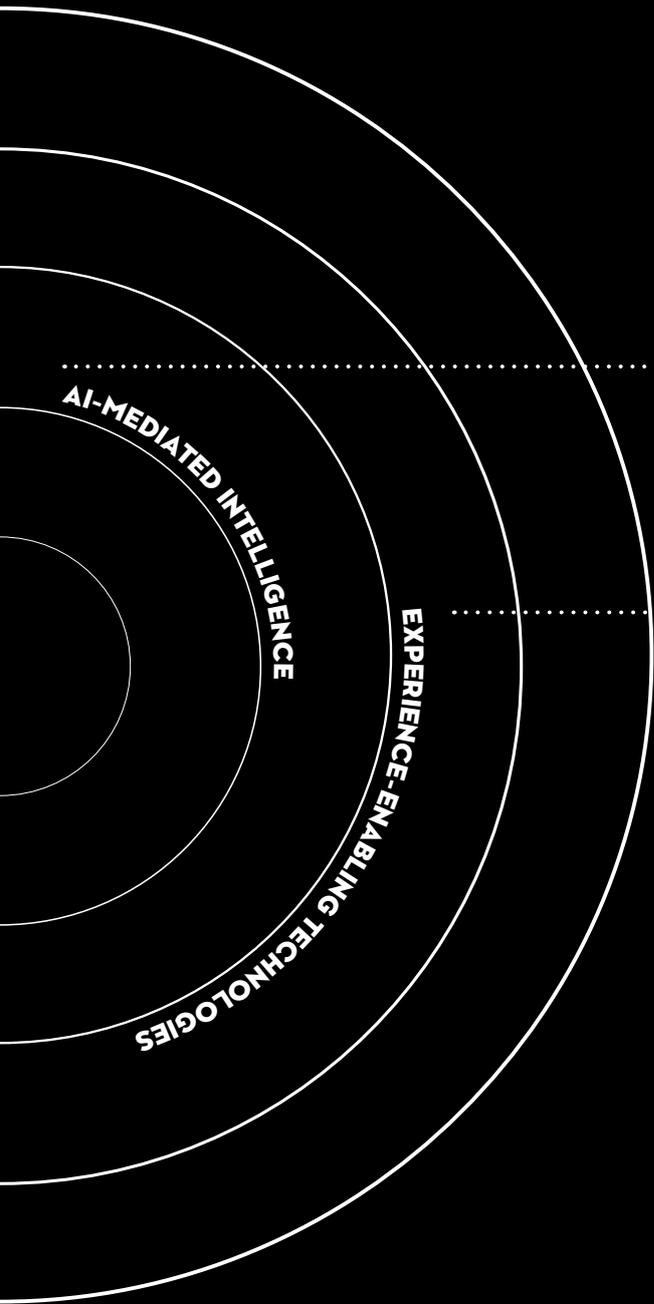
- Narratives, cultural norms and values, ethics, and governance.



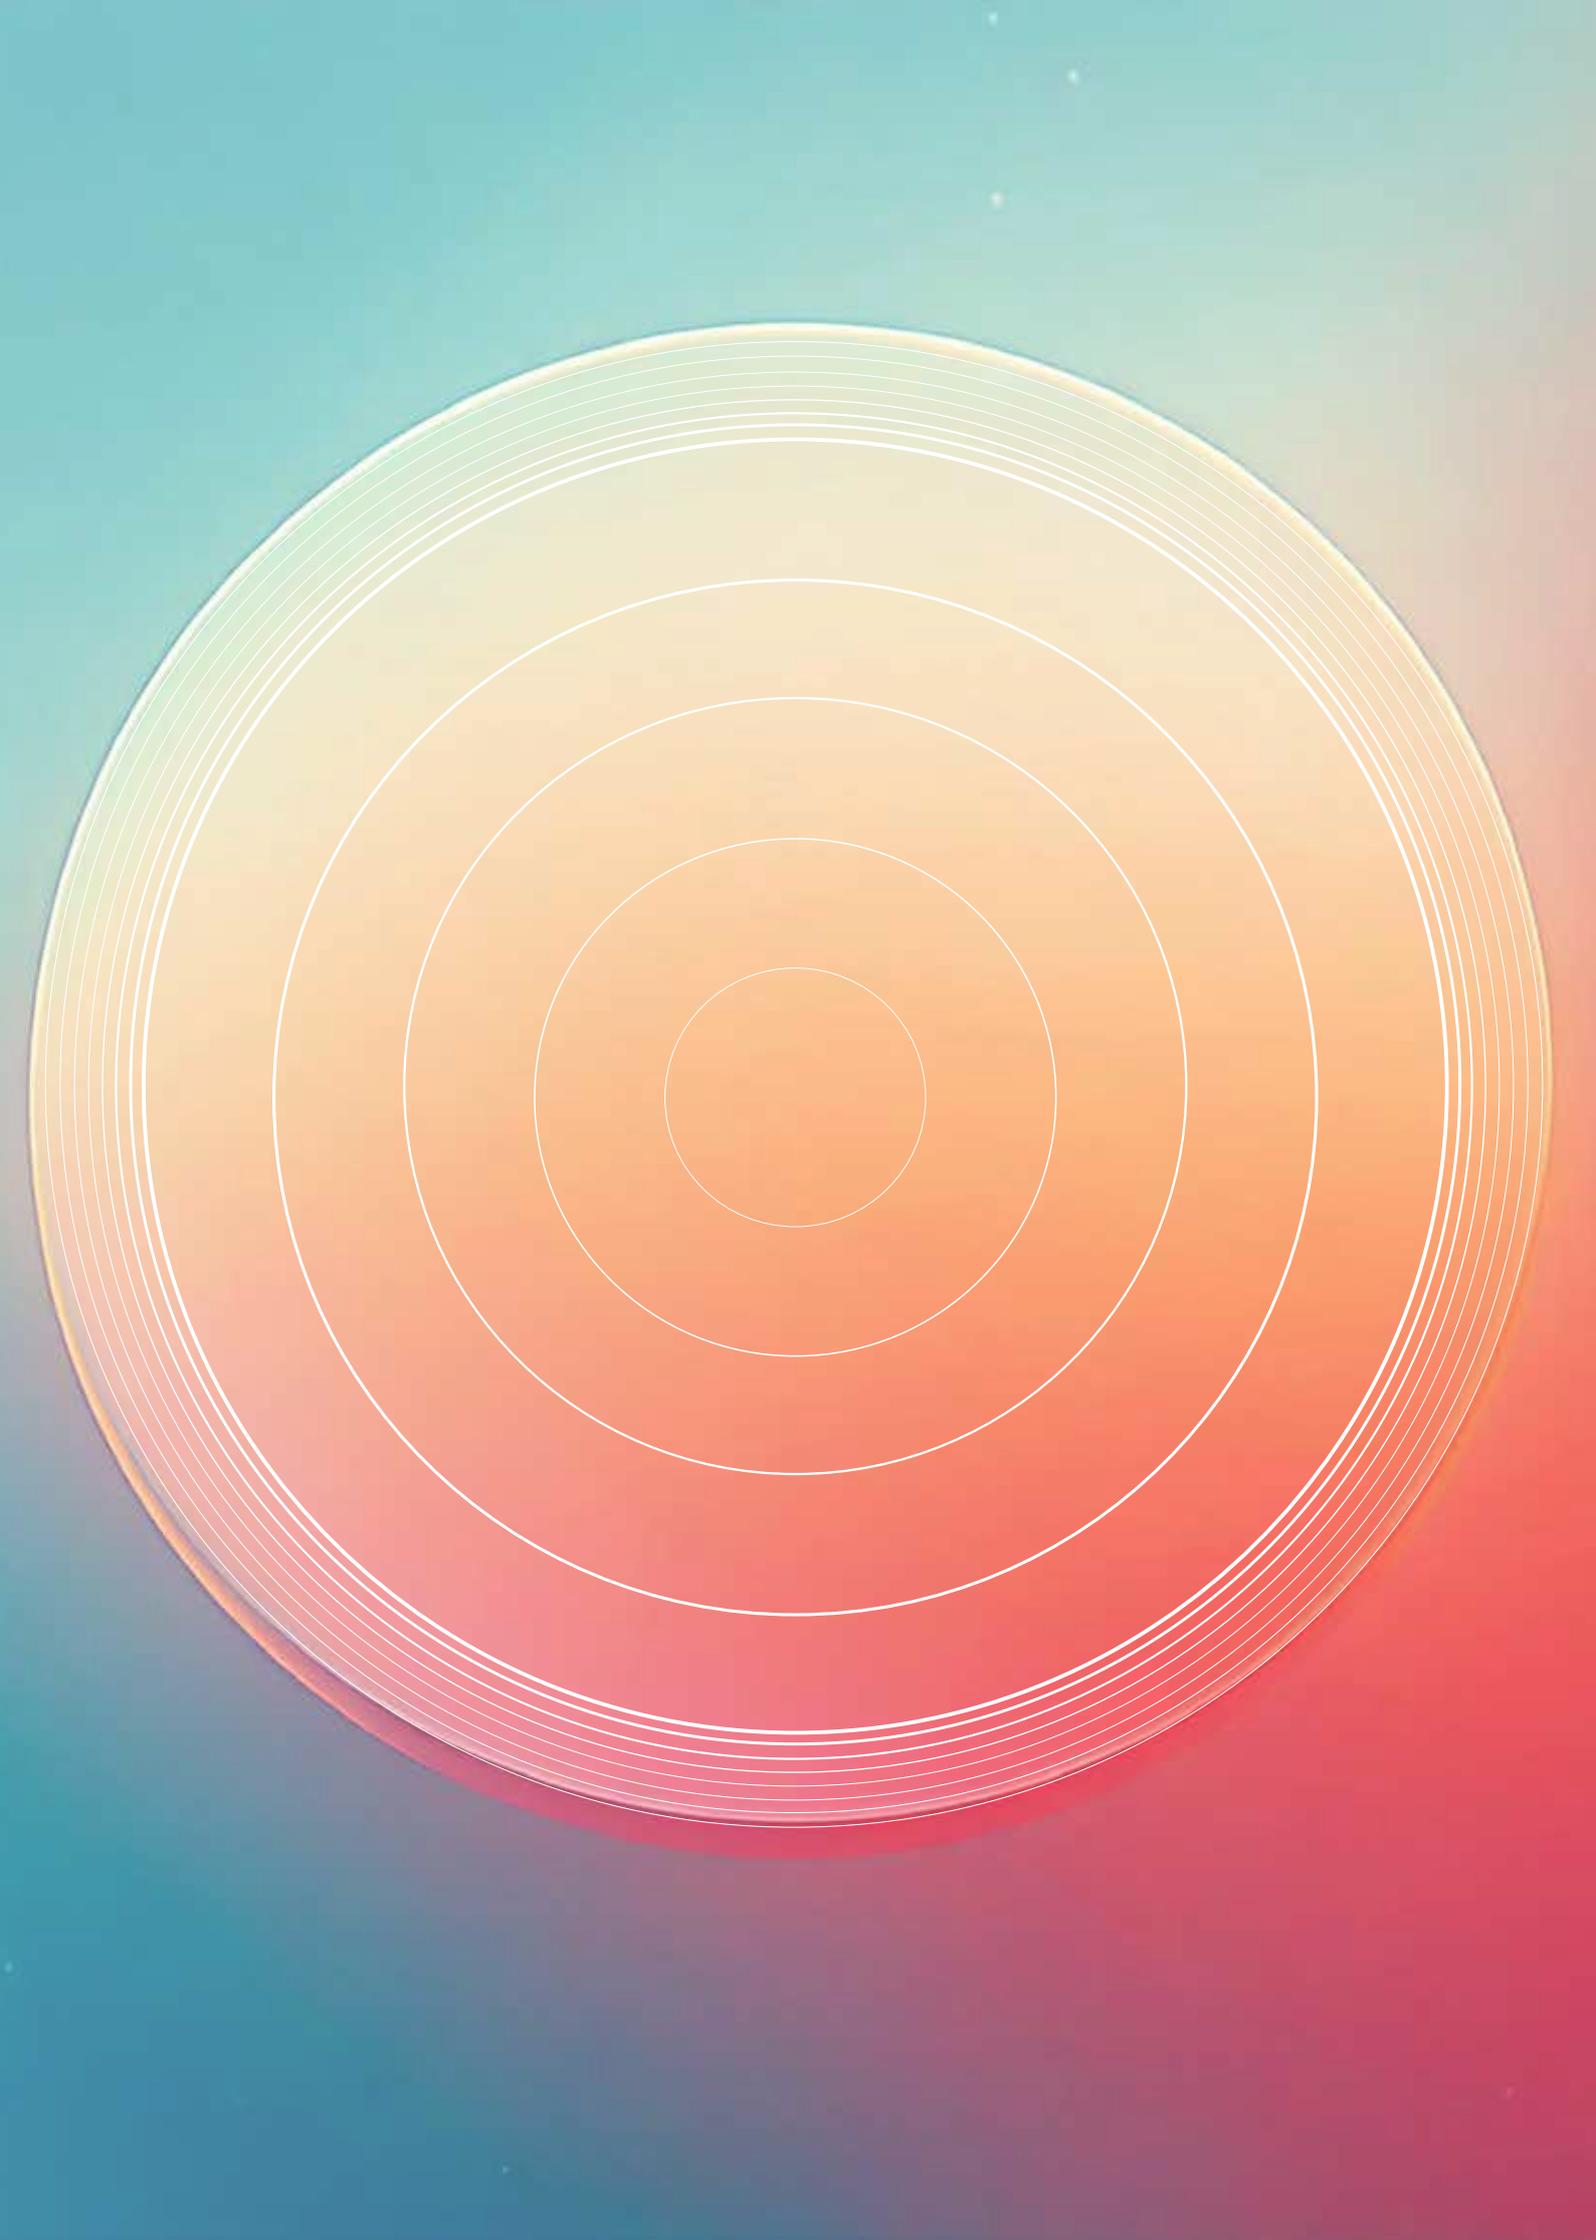
- Liquid identity, multiple representations, participatory.

- Scalable systems enabling everything in the inner layers.

- Includes interoperable platforms, quantum computing, and digital economies.



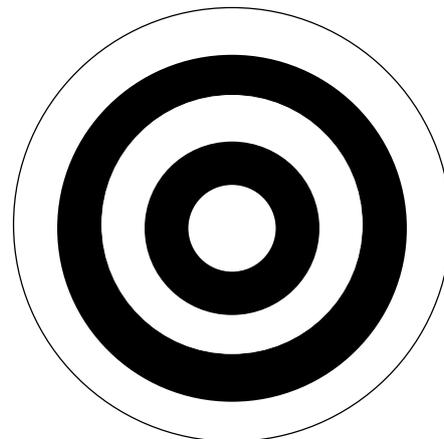
- AI systems personalising and adapting interactions.
- **Key enablers:** Contextual data and Agentic AI.
- Tools bridging digital and physical worlds.
- **Examples:** Immersive interfaces (XR), smart ecosystems, Physical AI.



10 KEY TRENDS & ISSUES TO WATCH IN 2025

The goal of this report is to examine the interconnected forces shaping emerging technologies, while also addressing some of the key trends and issues expected to define 2025. They are all part of the AI-mediated framework in different ways and look at the broader implications for individuals, businesses, and societies. We aim to highlight critical questions and encourage a future-driven understanding of the challenges and opportunities they bring. By addressing these issues holistically, this report seeks to provide a foundation for taking strategic decisions that reflect your preferred future and prepare you for what lies beyond.

1: The complementary power of big data and small data



The staggering growth of data generation shows no sign of slowing. An estimated 90% of the world's data has been created in just the past two years. Since 2010, the volume of data produced annually has surged year after year, reshaping how industries, organisations, and individuals interact with information. This explosion of data presents enormous potential but also introduces challenges, from information overload to polarisation driven by selective consumption and biases to the issues surrounding intellectual property (IP) conflicts and regulatory uncertainty, as seen in Generative AI.

In this landscape, the role of data in AI systems is expanding and making the interdependence of humans and AI even more refined. The two distinct but complementary directions for data are large, versatile models vs. smaller, task-specific solutions. It's getting increasingly important to understand the different types of models being developed.

Large Language Models (LLMs) have evolved from basic content generation to more complex reasoning, decision-making, and problem-solving mechanisms. They integrate vast datasets and diverse information, adapting to a wide range of needs. Even though there are doubts on how much better these big models can get, new reasoning and thinking models further improve LLMs, enabling multi-step logic and strategic problem-solving for real-world challenges. Innovations like Large Action Models (LAMs) enhance this by breaking down complex tasks and enabling real-time decisions. The debate over whether this development will eventually lead to Artificial General Intelligence (AGI) or Super Intelligence remains a contentious topic.

Meanwhile, Smaller Language Models (SLMs) prioritise precision and efficiency, making them optimised for more focused, task-specific applications. Open-source AI models are also paving the way for customised solutions, helping to accelerate progress by narrowing efforts toward products that can be applied effectively in the near term.

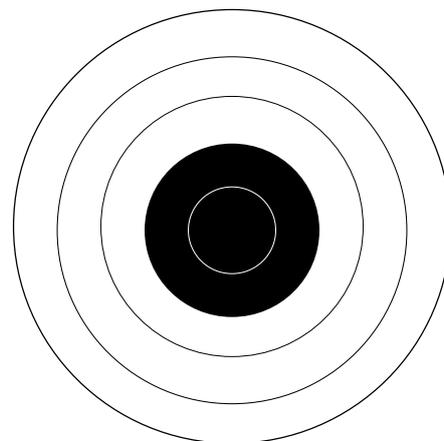
Advancements like Retrieval-Augmented Generation (RAG) reflect the increasing emphasis on combining static datasets with dynamic, real-time information – and gaining more control over the output. RAG systems enhance relevance and reliability by retrieving up-to-date knowledge tailored to specific contexts. The concept extends further with Agentic RAG, a system where autonomous agents dynamically retrieve, process, and act on data. This approach bridges the gap between information retrieval and execution, enabling AI to not only generate outputs but also reason through and act on them.

Together, these advancements highlight AI's growing role in tackling complex tasks across varied contexts and reflect a broader shift toward making data more usable, controllable and trustworthy. As systems grow in complexity, the focus is also on ensuring that AI serves specific, actionable purposes. By aligning models with human needs – whether through hyper-personalised services, intuitive interactions, or task-oriented applications – AI is becoming an indispensable mediator in how we engage with the digital world.

5 critical questions for companies and industries

1. How can we leverage both large and small AI models to optimise our operations and meet specific business needs while maintaining flexibility for future innovation?
2. What strategies can we implement to ensure our AI systems integrate static and real-time data in ways that enhance transparency, reliability, and ethical compliance?
3. How can we maintain control over AI outputs to ensure they align with our goals and values, while reducing risks tied to unintended or inaccurate results?
4. What safeguards can we adopt to mitigate risks like bias, misinformation, systemic failures, and a lack of control in AI-driven decision-making processes?
5. How can our organisation prepare for a future where hyper-specialised AI models increasingly dominate, while still taking advantage of breakthroughs in versatile, large-scale systems?

2: When context becomes king



As AI moves closer to our physical reality and our bodies, the context of our actions and interactions becomes increasingly vital in an AI-mediated world. Tacit knowledge – shaped by cultural norms, social cues, and bodily awareness – plays an essential role in how we act. In 2025, AI systems are progressing toward understanding and acting on this nuanced, contextual information, marking a significant step towards their ability to personalise experiences and adapt to unstructured, real-world settings.

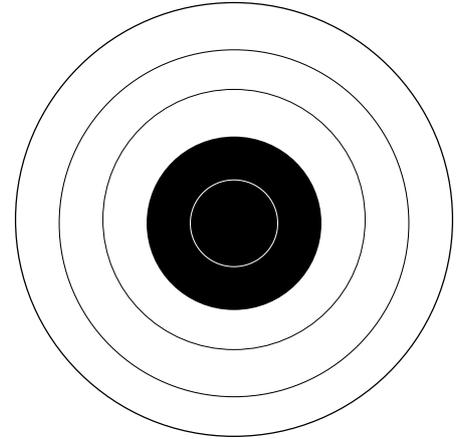
AI is entering a new phase where it can process not only structured and explicit data but also the tacit, contextual knowledge that underpins much of human experience. This tacit knowledge – rooted in cultural norms, social practices, and bodily awareness – is difficult to articulate or codify, yet it shapes how we navigate and understand the world. Technologies such as Google’s Project Astra are pioneering efforts to enable AI systems to interpret and act on such nuanced, contextual information. Project Astra focuses on developing systems that can perceive and process ambient signals, unspoken cues, and implicit behaviours. These advancements highlight how AI can extend beyond static, rule-based systems to provide real-time adaptability and deeper understanding of environments. For example, Astra’s work shows potential in creating AI solutions that can respond to complex, unstructured settings like a collaborative workspace or a dynamic healthcare environment.

As contextual AI evolves, it offers profound opportunities to personalise user experiences, optimise decision-making processes, and drive innovation across sectors. However, this potential is coupled with significant challenges. Contextual AI requires vast amounts of data to function effectively, raising critical concerns about privacy, data ownership, and the possibility of misinterpretation. Misaligned or biased interpretations of tacit knowledge could result in unintended consequences, especially in sensitive domains like healthcare, education, or legal systems.

5 critical questions for companies and industries

1. How might the adoption of contextual AI shift organisational practices, societal norms, and competitive dynamics – and what can businesses do to position themselves for these changes?
2. What new business models can emerge from AI's ability to adapt to unstructured environments and understand real-time context?
3. How can we strike a balance between leveraging contextual AI for personalisation and maintaining user trust and privacy?
4. What frameworks should be implemented to ensure that contextual data is collected, processed, and used transparently and responsibly while avoiding biases and ethical risks?
5. What safeguards can be put in place to minimise the risk of errors or misinterpretation in high-stakes environments?

3: Agentic AI - much more than chatbots



Agentic AI – also known as AI Agents - is one of 2025's most transformative advancements, marking a shift from passive tools to proactive systems the AI-mediated framework capable of reasoning, acting autonomously, and achieving complex goals. Moving from static to dynamic experiences. These AI agents are no longer limited to generating outputs but function as collaborators that orchestrate tools, external data, and reasoning frameworks to perform intricate tasks on behalf of humans. Their ability to adapt to dynamic environments and refine their performance over time makes them a game-changer across industries, from logistics to healthcare to personalised consumer applications.

At the core of agentic AI is its ability to break down complex objectives into manageable steps, allowing it to operate effectively in dynamic and unpredictable environments. This autonomy is complemented by a capacity to adapt to new circumstances, ensuring that the system remains aligned with its goals even as conditions evolve. Through iterative learning processes, agentic AI systems refine their performance over time, achieving increasingly optimised outcomes with minimal intervention.

Technological advancements such as cognitive architectures provide the underlying capabilities for these systems to reason, solve problems, and navigate uncertainty. Additionally, the integration of tools and external data sources enhances their versatility, enabling them to access specialised information or execute tasks that require domain-specific knowledge.

Agentic AI is already finding applications across industries. In logistics, these systems optimise supply chains autonomously, reducing inefficiencies and responding to real-time changes. In healthcare, they assist with diagnostics, treatment planning, and even patient care. Consumer-facing uses, such as personal AI assistants or education-focused AI companions, illustrate the potential for agentic systems to integrate seamlessly into everyday life, offering enhanced support and personalised experiences.

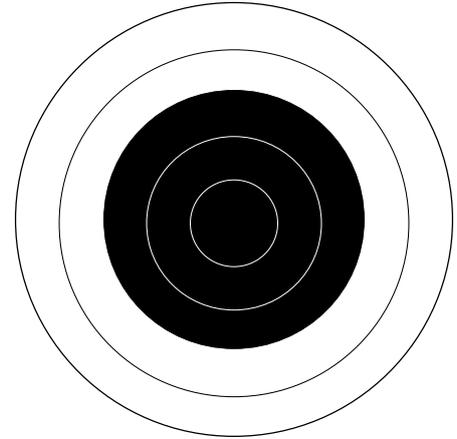
However, as these systems become more sophisticated and human-like in their behaviours, there's a growing risk of anthropomorphisation – attributing human emotions, intentions, or consciousness to AI. This can lead to misplaced trust or unrealistic expectations, complicating how these systems are understood and relied upon. It's crucial for businesses and users to remain mindful of this risk, ensuring that agentic AI is deployed transparently and understood for what it is: a powerful tool, not a substitute for human judgment or empathy.

While the potential benefits of agentic AI are immense, its autonomous nature introduces significant challenges. The capacity to act independently raises concerns about accountability, especially when actions lead to unintended consequences. Ensuring that these systems perform, and in the same time act ethically, transparently, and reliably. The development of governance frameworks that address these concerns will be essential to unlocking the full potential of agentic AI while mitigating its risks.

5 critical questions for companies and industries

1. How can Agentic AI be used to optimise workflows and decision-making processes in ways that directly enhance operational efficiency and benefits the individual?
2. What strategies can ensure autonomous AI agents optimise operations while minimising risks to brand reputation and customer trust?
3. What investments in infrastructure, talent, and partnerships are required to fully realise the potential of Agentic AI in achieving business goals?
4. What safeguards are needed to ensure the accountability, transparency, and safety of autonomous AI systems in high-stakes applications?
5. How should we prepare for and adapt to new regulatory policies designed to govern the deployment of agentic AI, particularly in areas like liability and compliance?

4: AI Clones & our own multiple representations



By 2025, AI clones – lifelike digital representations of individuals – may begin to see actual practical applications in isolated use cases. These clones that can replicate your appearance, voice, and even decision-making style, are enabling new possibilities for both personal and commercial use if the models have enough context. AI clones represent a significant shift in how individuals and businesses interact with technology. The possibility to create scalable content and deliver consistent experiences is straight forward. From streamlining and personalising operations to rethinking engagement.

AI clones are advancing rapidly and are already being used to generating highly realistic influencer-style videos for platforms like YouTube and TikTok. These videos can make makeup tutorials, product reviews, and travel vlogs, all created by AI-driven personas that look and sound like real people – and might or might not be a clone of a real person. This development can offer businesses a cost-efficient way to scale content creation and customer engagement without much human involvement.

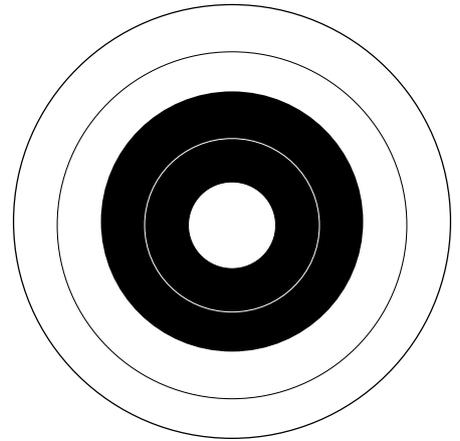
Beyond marketing, the potential of AI clones extends to areas such as customer service, training, and even executive representation. Imagine an AI clone of a senior leader delivering structured insights in a client meeting, ensuring consistency and professionalism while freeing up time for strategic priorities. Or consider an AI clone of a subject-matter expert who is available 24/7 to provide consistent, personalised training sessions, available globally and tailored to individual learning needs.

However, this technology also presents challenges. As AI blurs the line between human and synthetic content, questions emerge about how it impacts our perception of authenticity and the trustworthiness of interactions. In some cases, consulting an AI clone may even be preferred to the individual themselves, especially in contexts requiring structured, data-driven insights.

5 critical questions for companies and industries

1. What frameworks are needed to address intellectual property rights, ownership, and ethical considerations for AI clones, ensuring individuals and organisations retain control and protection?
2. What frameworks and safeguards are needed to ensure AI clones remain accurate and current representations of individuals, align with personal or organisational values?
3. In what scenarios might AI clones outperform humans in providing insights or services, and how can organisations manage this transition transparently while addressing professional liability and team dynamics?
4. How can organisations balance the use of AI clones for hyper-personalised interactions while fostering genuine human relationships, maintaining organisational culture, and addressing potential effects on employment and union rights?
5. What strategies can ensure equitable access to AI cloning technology, prevent its benefits from being concentrated among the privileged, and prepare for potential technical failures or performance evaluations against human benchmarks?

5: Bringing intelligence to the physical world with Physical AI



In 2025, the evolution of Physical AI is accelerating, bridging the gap between digital intelligence and real-world interaction. Physical AI represents a transformative step in the evolution of artificial intelligence, where AI systems fuse real-world sensor data to enable real-time perception, understanding, and reasoning about the physical world. Unlike traditional AI applications that operate in purely digital environments, Physical AI is embedded within physical systems, leveraging sensor inputs to interact with and adapt to dynamic, real-world conditions in an AI-mediated world.

Current advancements such as Large Behaviour Models (LBMs) and Vision-Language-Action (VLA) models are setting the stage for this transition – just to mention a few. LBMs aim to predict and align with human behaviour, making AI agents better at anticipating needs and solving multi-step problems. VLA models extend this capability by linking visual recognition, language comprehension, and physical action, making these agents more versatile and effective in executing tasks.

This integration is powered by advancements on the Internet of Things (IoT), where networks of interconnected devices collect vast amounts of data from the physical environment. Everyday examples, such as robotic vacuum cleaners navigating autonomously or smart thermostats optimising home temperatures, showcase how Physical AI is already part of our lives. However, the potential extends far beyond these applications. The vision for Physical AI includes encoding the entire physical world, uncovering hidden patterns and structures in behaviours to revolutionise industries such as manufacturing, healthcare, and urban planning.

By bridging the gap between digital intelligence and the physical world, Physical AI promises to enhance efficiency, enable adaptive automation, and open new possibilities for real-time decision-making. Yet, this shift also brings critical challenges. The collection and use of physical data raise concerns about privacy and ethical governance, while the rapid advancement of Physical AI risks outpacing the development of regulatory frameworks. As these technologies continue to expand, their societal and ethical implications will demand careful consideration.

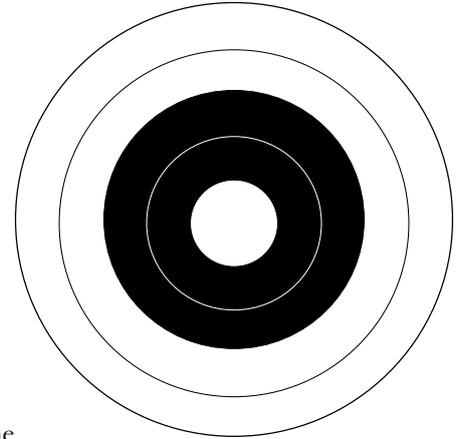
There is an overlap between physical AI systems and agentic AI (like autonomous robots), but not all agentic AI systems have physical embodiment (like software agents). The key difference is that agentic AI is defined by autonomous decision-making and goal-directed behaviour, while physical AI is defined by its connection to physical hardware and real-world interaction.

As Physical AI becomes more embedded in everyday life, 2025 represents a critical time for companies to understand its potential while navigating its complexities.

5 critical questions for companies and industries

1. How can we leverage Physical AI to improve operational efficiency and deliver adaptive, real-time solutions?
2. What steps can we take to integrate Physical AI seamlessly into existing workflows, balancing innovation with ethical considerations and societal expectations?
3. How can we deploy Physical AI systems to maximise value while maintaining accountability for their actions in complex, real-world environments?
4. What strategies can we implement to secure and responsibly manage sensor-driven data, protecting customer privacy and maintaining trust?
5. How can we ensure our use of Physical AI complies with emerging regulations while proactively addressing concerns about autonomy and decision-making in physical spaces?

6: The rise of humanoid robotics



The development of humanoid robotics represents a transformative milestone in the evolution of AI, where advanced Generative AI and computer vision converge to create machines capable of mimicking human form, function, and behaviour. These robots are envisioned to perform a wide range of tasks, from industrial operations to personal assistance, while seamlessly integrating into environments designed for humans. This shift is part of a broader trend in robotics aimed at creating adaptive, intelligent systems that interact naturally with their surroundings.

According to NVIDIA, humanoid robotics can be classified into three categories. 1) Agentic AI (Digital Robots): Digital-only robots, such as virtual assistants or avatars, that operate entirely in the digital realm. 2) Autonomous Units: Robots designed to perform specialised tasks in physical environments, like self-driving cars or drones. 3) General-Purpose Robotics (Humanoids): Robots modelled after the human form, capable of performing a wide array of tasks and adapting to dynamic, unstructured environments.

Recent advancements highlight the promise of humanoid robotics. These machines leverage Generative AI to process and generate data in real time, enabling adaptive learning and decision-making. At the same time, computer vision allows them to interpret their surroundings, recognise objects, and navigate complex spaces. Unlike earlier robots focused on narrow tasks, humanoids are becoming multi-functional, with the potential to serve in roles requiring both physical dexterity and cognitive reasoning. This shift is supported by significant investments from major players like Microsoft and OpenAI, accelerating the transition of humanoid robotics from research labs to real-world deployment.

Technological advancements in hardware have also played a crucial role. Cutting-edge humanoid robots now incorporate lifelike designs and intricate functionality inspired by human biology. For instance, some utilise synthetic muscle fibres and sensory systems, allowing them to emulate natural movement and respond to physical stimuli. These developments are underpinned by foundational technologies such as visuomotor models, which combine visual perception with motor actions, and AI systems capable of learning complex behaviours.

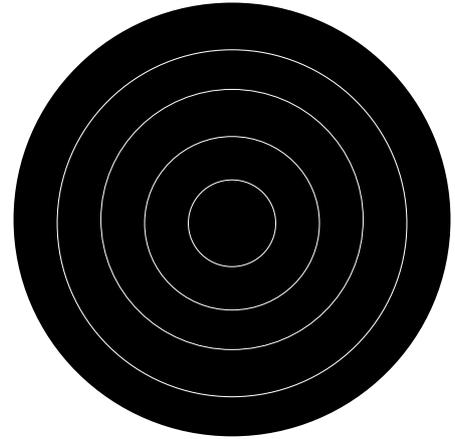
The rise of humanoid robotics offers transformative potential across industries. In manufacturing, humanoid robots can operate in highly automated production lines, complementing human workers by performing repetitive or hazardous tasks. In healthcare, they could assist with caregiving, rehabilitation, and even surgeries. In homes, humanoid robots could act as personalised assistants or companions.

Despite their promise, humanoid robotics also raises critical ethical and societal questions. The introduction of human-like machines into workplaces, homes, and public spaces prompts concerns about labour displacement, privacy, and the psychological effects of interacting with machines designed to emulate human behaviour. Questions about accountability and the potential misuse of humanoid robots further complicate their adoption. Ensuring that these systems are developed responsibly, aligned with societal values, and governed by transparent policies will be crucial as humanoid robotics becomes an integral part of our future.

5 critical questions for companies and industries

1. How can we determine the right time to invest in humanoid robotics, balancing their technological maturity with our organisational readiness?
2. How might humanoid robotics reshape traditional concepts of labour, and what can we do to support workers in adapting to these shifts?
3. How can humanoid robotics help address labour shortages in industries struggling to attract or retain workers or be used to improve workplace safety, particularly in roles that involve repetitive tasks or hazardous environments?
4. What ethical frameworks are needed to address the societal and psychological impacts of creating robots that closely mimic human form and behaviour?
5. What long-term societal shifts might arise from the integration of humanoid robots, and how can we proactively address their potential consequences?

7: The rise of dual-use technology



In 2025, the development of dual-use technologies is accelerating, blurring the lines between civilian and military applications. Dual-use technologies are innovations that can be applied in both civilian and military contexts. Advances in artificial intelligence, drones, quantum computing, and other fields are driving innovation across sectors while simultaneously raising ethical and regulatory challenges in defence and security contexts. As geopolitical tensions rise, the potential for misuse and unintended consequences underscores the urgent need for governance frameworks.

While this overlap drives innovation, it also introduces complexities around ethical considerations and regulatory frameworks. Civilian technologies can unintentionally become tools for surveillance or conflict, raising questions about accountability and unintended consequences. The expansion of dual-use technologies into areas such as autonomous weapons, AI-enabled reconnaissance, and offensive cyber capabilities heightens these concerns, particularly seen in the light of increasing geopolitical tensions.

As the pace of dual-use technology development accelerates, new risks emerge. Civilian infrastructure may be co-opted for military purposes, while a lack of clear governance frameworks increases the potential for misuse. The line between civilian and military applications grows increasingly blurred, requiring global collaboration to address ethical dilemmas, regulate technology transfer, and ensure accountability in development and deployment. Without such frameworks, the proliferation of dual-use technologies could undermine trust, destabilise international relations, and erode safeguards intended to protect civilians from harm.

5 critical questions for companies and industries

1. How might the rapid advancement of dual-use technologies change the way businesses collaborate with governments or other organisations?
2. How can we ensure our dual-use innovations create value for civilian applications while minimising the risks of exacerbating geopolitical tensions?
3. How can we ensure that user data collected through civilian technologies is safeguarded from being repurposed or deployed in military applications without explicit consent or oversight?
4. What measures can we take to build trust with stakeholders, demonstrating transparency and accountability in the development of dual-use technologies?
5. What ethical considerations should guide our decision-making when developing technologies with both civilian and military potential?

8: The uncertainty of scaling



Scaling is a defining challenge for 2025, as it determines how effectively emerging technologies can transition and the possible impact on the ecosystem it might have. Whether it is computing power, productivity, investment, organisational growth, or innovation, the ability to scale shapes the trajectory of technological and economic progress. However, scaling systems often encounter thresholds where growth introduces complexities that can either accelerate breakthroughs or create significant obstacles. While scaling in these areas promises unprecedented advances, it also introduces critical uncertainties.

The scale of compute is particularly crucial for technologies like AI and quantum computing. While increased computational power enables sophisticated models and unprecedented advances, it also brings challenges such as energy consumption, infrastructure bottlenecks, and diminishing returns. For example, next-gen GPUs and 6G infrastructure are critical for real-time AI applications, yet they require significant resource investments.

Similarly, the scale of productivity reflects how technological advancements translate into real-world efficiency. AI, robotics, and automation offer immense promise, but their potential hinges on complementary investments in skills, organisational change, and infrastructure. Without alignment, gains may remain uneven, benefiting certain sectors or regions disproportionately.

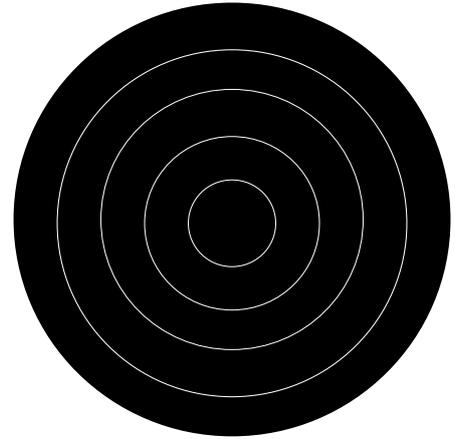
The scale of investment highlights the financial resources required to sustain transformative growth. While investments in e.g. AI and quantum computing are at record highs, disparities in access and the risk of speculative bubbles and over-concentration may threaten long-term stability. Sustaining these investments without destabilising markets or reinforcing inequalities is essential.

Finally, scaling innovation requires balancing incremental progress with radical breakthroughs. This drives the entire ecosystem, pushing the boundaries of what is possible. However, innovation does not always scale linearly. Larger organisations and systems may struggle to maintain agility and creativity, potentially stalling innovation as complexity increases. Why radical innovation will be more important than ever. Balancing incremental and radical innovation with governance and resilience will be crucial as systems scale further.

5 critical questions for companies and industries

1. How can we manage the scale of compute to drive technological breakthroughs while addressing energy consumption and infrastructure limitations?
2. How can we maximise productivity gains from emerging technologies within our organisation, ensuring benefits are effectively realised across teams, processes, and operations?
3. How do we balance the scale of investment to sustain growth while avoiding speculative bubbles or over-concentration in specific sectors?
4. How can we foster both incremental and radical innovation within our organisation to maintain responsiveness, creativity and relevance as we scale?
5. What should be our priority investments – in hardware, organisational development, or radical innovation – to ensure sustainable scaling without overextending resources?

9: Environmental implications of an AI-mediated infrastructure



As AI becomes increasingly central to technological ecosystems, its environmental impact demands urgent attention. Can we justify the energy demands of AI-mediated systems while maintaining a commitment to sustainability?

The rapid growth of large-scale AI models, real-time data processing, and autonomous systems has escalated energy consumption, raising critical concerns about ecological responsibility. By 2025, these issues are becoming more pressing as AI expands into nearly every domain, creating opportunities for sustainability innovation while also amplifying risks of environmental stress. For individual companies, navigating these challenges involves carefully considering the types of AI models employed and their associated energy demands. Not all tasks require resource-intensive systems; simpler classification tasks, for instance, may be effectively handled by lightweight models or even traditional machine learning approaches. By matching the complexity of the model to the complexity of the problem, organisations can optimise performance while minimising energy use.

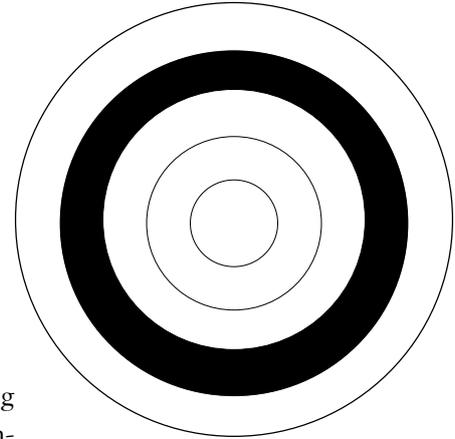
Larger companies, particularly those reliant on complex, multi-territory operations, should also assess how different infrastructure choices affect scalability and sustainability. Understanding how smaller, decentralised models – such as those utilising edge computing - compare to centralised, resource-heavy systems can guide decisions on whether to build lighter, more efficient AI solutions or invest in robust centralised infrastructures for specific needs.

As AI-mediated systems continue to evolve, there is also growing interest in designing more energy-efficient AI models, minimising redundant computations, and leveraging AI to optimise energy use in other industries. These developments highlight a dual role for AI: as both a potential driver of sustainability and a contributor to environmental challenges. By staying attuned to these developments and fostering awareness of their ecological dimensions, we can better understand how AI fits within a sustainable future. This is not about fixed solutions but about maintaining a thoughtful approach as these technologies evolve and reshape the world.

5 critical questions for companies and industries

1. How can we assess and actively reduce the environmental footprint of their AI-mediated systems, from data centres to edge devices?
2. What opportunities exist to integrate energy-efficient computing, such as edge computing or more sustainable chips, into AI infrastructures to align with environmental goals?
3. What role can we play in driving innovation for energy-efficient AI models and minimising redundant computations in our systems?
4. What collaborative opportunities exist between businesses, governments, and regulators to accelerate the development of policies addressing the environmental impacts of AI?
5. How can businesses communicate their environmental efforts around AI systems transparently, building trust with consumers and stakeholders?

10: Quantum coming closer



Quantum computing, once considered a distant breakthrough, is moving closer to practical use. With advancements in hardware, algorithms, and quantum infrastructure, we are beginning to see the first tangible steps toward its integration into industries and research. Quantum computers leverage the principles of quantum mechanics to process information in ways that classical computers cannot, making them capable of solving problems that were previously intractable.

Unlike classical computers, which rely on binary states (0s and 1s), quantum systems use qubits that can exist in multiple states simultaneously, thanks to phenomena like superposition and entanglement. This allows quantum computers to perform parallel computations at an unprecedented scale. Potential applications span a range of fields, including cryptography, materials science, drug discovery, logistics optimisation, and climate modelling. For example, quantum algorithms could revolutionise supply chains by identifying optimal routes in seconds or accelerate the discovery of new pharmaceuticals by simulating molecular interactions at an atomic level.

Concrete advancements underscore this progress. Google has introduced Willow, a quantum processor capable of solving tasks in five minutes that would take current supercomputers an estimated 10 septillion years. This breakthrough demonstrates remarkable improvements in quantum error correction and computational capabilities.

However, quantum's arrival also comes with significant challenges. One major concern is its potential to break existing encryption protocols, which would render many current cybersecurity systems obsolete. This has led to a race for "quantum-safe" encryption methods to protect sensitive data in the post-quantum era. Additionally, the hardware required for quantum computing remains fragile and expensive, with scalability still a major hurdle.

As quantum technology edges closer to mainstream use, its societal and geopolitical implications must be carefully considered. The ability to harness quantum power could create disparities between nations and organisations with access to the technology and those without. This makes collaboration and the development of ethical frameworks essential to ensuring that quantum advancements benefit humanity as a whole.

5 critical questions for companies and industries

1. How can we balance investments in quantum technology with the uncertainties surrounding its maturity and scalability?
2. What specific applications of quantum computing should we prioritise to gain a competitive advantage in areas like optimisation, materials development, or data analysis?
3. How can we prepare our cybersecurity systems and infrastructure to address the potential risks posed by quantum computing, such as the obsolescence of traditional encryption?
4. What strategies can we adopt to ensure equitable access to quantum advancements, avoiding monopolisation by a few dominant players and fostering collaboration?
5. How can we – and should we – work with governments and other stakeholders to shape policies that address ethical considerations, geopolitical impacts, and the societal implications of quantum computing?

Rethinking foresight in an AI-mediated world

We often struggle to imagine the broader implications of transformative technologies, which limits our ability to prepare for radical change. This lack of creative imagination can arise from a reliance on linear thinking, a fear of the unknown, or a tendency to frame the future through the lens of present-day challenges. As a result, technological progress is often viewed narrowly, focusing on incremental applications rather than transformative potential. For instance, while much attention is given to how AI might automate tasks, less consideration is afforded to its capacity to reshape societal structures or re-define human agency.

The trends and issues discussed above, along with the critical questions they raise, are deeply interconnected and constantly evolving. While they all converge toward a more AI-mediated future, a range of possible outcomes must be considered. By challenging assumptions and exploring “what if” scenarios, we can identify risks and opportunities before they fully materialise. For example, envisioning a future where deepfakes are indistinguishable from reality can guide the development of verification tools or policies to protect trust in digital environments. Similarly, speculative design can simulate AI-driven systems to explore their societal implications, from privacy concerns to ethical dilemmas.

A deeper exploration of the potential scenarios that arise from the trends and issues discussed falls outside the scope of this report. Yet one final and equally crucial question to consider is how futures thinking and foresight itself can benefit from AI. One answer to the question of how organisations can work with the future in an AI-mediated world is that it requires an equally AI-mediated foresight toolbox.

A foresight toolbox that combines AI's capabilities allows foresight practitioners to sift through vast data sets, identify weak signals, and simulate complex systems, enabling faster, deeper insights. When paired with human-led creativity and methodologies like strategic foresight and horizon scanning, it can deliver foresight that is more dynamic and adaptive.

In the space where AI, cultural insight, and futures thinking converge, a new kind of strategist emerges – part technologist, part anthropologist, and part data wizard, and uniquely equipped to uncover connections others miss. The future isn't only about disruption or replacement; it's also about opening new territories of possibility. The ability to synthesise across disciplines, technologies, and human experiences to envision and prepare for what's possible next will be a powerful and important tool in an increasingly AI-mediated world.

This future is also not a fixed trajectory but a constantly evolving landscape, shaped by the collective choices we make today. Emerging technologies are not isolated innovations or fleeting trends in the hype cycle. They are deeply interconnected, forming the building blocks of an increasingly complex technological ecosystem that influences how we live, work, and interact. These systems, which bridge digital and physical realms, hold the potential to transform industries, enhance human experiences, and drive meaningful progress – but only if we approach them with purpose and curiosity for the possible futures.

“One answer to the question of how organisations can work with the future in an AI-mediated world is that it requires an equally AI-mediated foresight toolbox.”

STRATEGIC FUTURE-DRIVEN QUESTIONS TO ASK IN 2025

If strategic foresight and a prepared mindset are the tools for navigating the complexities of an uncertain future, then the questions we ask become just as important as the answers we seek. In 2025, the pace of technological change, coupled with the convergence of emerging innovations, demands a proactive and inquisitive approach. Businesses, industries, and societies must consider not only where we are heading but also the choices and actions that will shape that trajectory.

The right questions can help uncover hidden connections, clarify uncertainties, and guide decision-making in a way that prepares us for future challenges and opportunities. They enable us to look beyond the immediate horizon, anticipate shifts in behaviour, and design strategies that balance innovation, trust, and sustainability. In an era where adaptation and foresight are critical, these questions serve as the foundation for developing a long-term perspective and ensuring relevance in an ever-changing landscape.

Here are some of the most critical questions to be aware of:

- **How can we understand the layers of emerging technologies within and around our industry to uncover their interdependencies and develop cohesive strategies for the future?**
- **What tools or practices can we adopt to identify the right moment to invest in technologies, balancing maturity, organisational readiness, and market dynamics?**
- **How do we foster a culture of resilience and adaptability to navigate rapid technological acceleration while staying aligned with long-term goals?**
- **What actions can we take today to ensure equitable access to emerging technologies and prevent the deepening of societal or economic divides?**
- **How can we prepare for and anticipate scenarios where technological disruptions challenge established business models or global systems?**

- **How do we rethink consumer engagement in an AI-mediated ecosystem, ensuring personalisation enhances trust and autonomy without compromising agency?**
- **What leadership capabilities are necessary to enable proactive, future-focused strategies that embrace uncertainty and foster sustainable growth?**
- **How can businesses balance the pursuit of innovation with the need for trust, transparency, and ethical practices across AI-driven ecosystems?**
- **How can we align incremental advancements with radical, future-driven innovation to maintain relevance and competitiveness?**
- **What strategic steps can we take to better visualise and communicate transformative futures, fostering a shared language for navigating change?**

CIFS' FORESIGHT FRAMEWORK

The future is inherently unpredictable, and many factors may come together in complex ways to create surprising futures in a non-linear world. There are no definitive answers about what the future will hold. Instead, we must start preparing for the possible futures and create a prepared mindset, implementing a long-term perspective. With futures thinking we are trying to learn about the development through the lens of the future, in contrast to, per default, trying to understand the future primarily from the perspective of the past and the present. It is a structured approach that allows organisations to consider alternative future outcomes representing novel perspectives and contexts. Hence, futures thinking can help organisations widen their strategic perspective and devise strategies that are more resilient across different futures.

Futures thinking adopts an 'outside-in' approach, focusing on potential changes in the external environment – including outcomes 'beyond the numbers' that might otherwise be overlooked. This, in turn, influences the strategic environment and consequently strategic decisions. The process is both explorative and deductive: it aims to derive an understanding of plausible and consistent futures through the close examination of driving forces and critical uncertainties. Our approach to futures thinking is anchored in co-creation and broad participation and follows a structured, multi-step process. The primary focus is on fostering futures literacy – the ability to employ anticipation and imagination for improved strategic decision-making, understanding key drivers of change, and exploring emerging opportunities.

ABOUT COPENHAGEN INSTITUTE FOR FUTURES STUDIES

Our purpose: To help people and organisations imagine, work with, and shape their future.

The Copenhagen Institute for Futures Studies is an independent, non-profit think tank established in 1969. By building the capabilities necessary to address potential futures, we help create a society fit to meet the challenges and grasp the opportunities we face.

We do this by applying our unique approach to futures studies and foresight, combined with more than 50 years of global experience and contributions to the field, working with organisations across the public, private, academic, and civic sectors, as well as with the general public.

Our vision is a futures literate world where everyone has the right and mandate to engage with the future, participate, and visualise change, so they can create the best possible future for themselves, society, and the planet.

Read more at **WWW.CIFS.DK**

**COPENHAGEN
INSTITUTE
FOR FUTURES
STUDIES**



