Looking Glass

Bringing tech-led business changes into focus

thoughtworks

Strategy. Design. Engineering.

2024

Introduction

Typical technology trend reports will identify what they think demand your attention. Looking Glass is different: it's designed to help you use your attention more effectively. It does this by offering you a way of looking at trends, a way of analyzing and interpreting emerging technologies so you can make informed decisions for your organization.

Making those decisions has arguably never been more challenging. Today's rapid pace of technological change can make it difficult to form a clear view on what's ahead and where you're likely to see the most value for your investments. While the shockwaves caused by significant technologies like ChatGPT will undeniably shape what the future looks like, the past certainly shouldn't be read as a manual for what's next. That's why we publish Looking Glass.

In this 2024 edition, we've identified over 100 trends through five lenses that we see as defining the future of technology in business. Some of these trends are already shaping the way organizations operate, while others are somewhere on the horizon, attracting attention and conversation but firmly rooted in the future. As a business leader, it's critical to take a high-level view on the key trends shaping the world's businesses and technologies — regardless of the trends happening now or further in the future. Looking Glass is a tool for doing just that — and the lenses on which it is built offer precisely this all-important high-level view. Ultimately, it ensures your organization is adaptable, resilient and well-primed to weather or leverage the technology changes that are an inevitable part of modern life.

—

Rachel Laycock
Chief Technology Officer, Thoughtworks

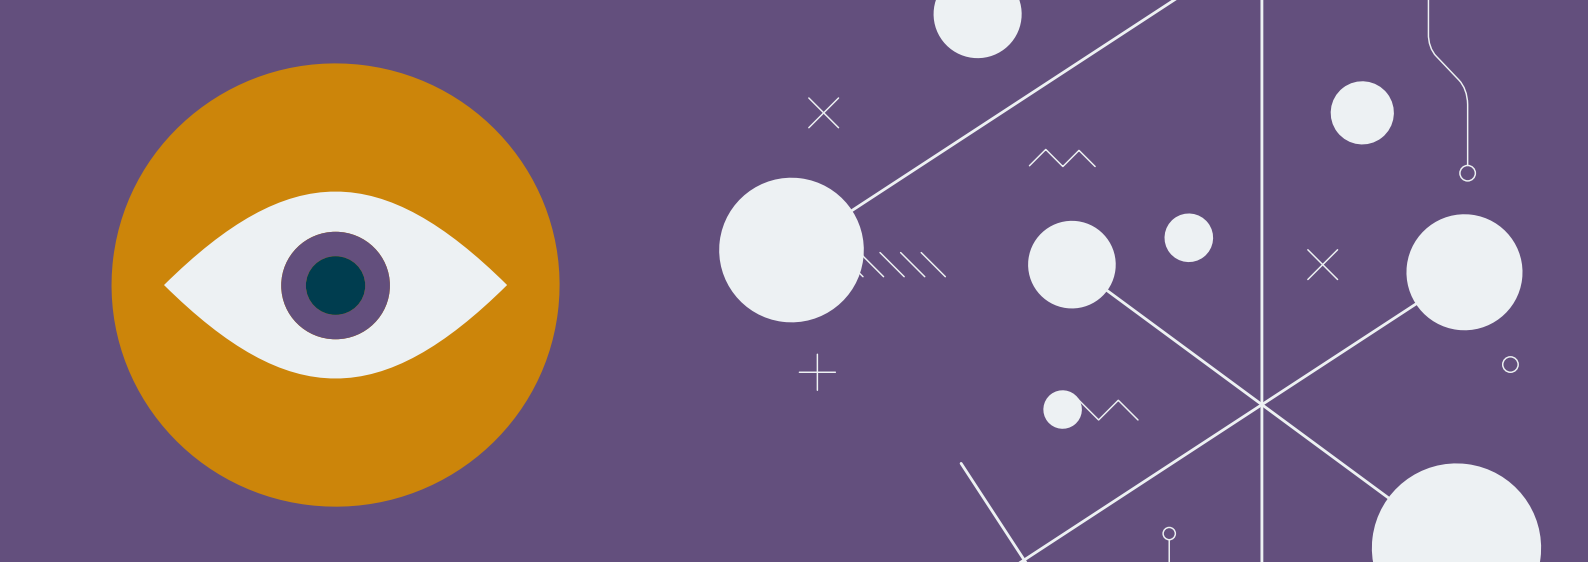

AI everywhere

Leveraging cutting-edge breakthroughs to scale your business

Generative AI (GenAI) has captured the spotlight, but in reality it's just one aspect of a much broader field that is advancing on multiple fronts. One of the things GenAI has proven is that AI can be made available, accessible and applicable to more people. This democratization has prompted a rush of experimentation and investments in everything from smartphone alternatives to startups working on the next ChatGPT.

What all this will mean for organizations at the day-to-day level is less clear. We are firm believers that AI is already and will continue to have a major impact on some processes integral to being a digital business, notably software development, enabling enterprises to build and bring products to market faster.

It's important to understand that, for all the excitement, AI won't always be a transformational force. Heavyweight commercial large language models (LLMs) are powerful but — at least for now — generally too expensive for most organizations to use for anything at scale. The buzz surrounding ChatGPT means that it can become a 'hammer looking for a nail,' with companies rushing to integrate it into processes when it may not actually be necessary, or the right tool for the job.

Organizations need to put a few fundamental building blocks in place before they can take advantage of the AI breakthroughs that seem to be emerging every day

Organizations need to put a few fundamental building blocks in place before they can take advantage of the AI breakthroughs that seem to be emerging every day. One is a solid data strategy, as outlined in our data platform lens, that ensures a base level of relevant, credible and traceable data is readily available to feed into AI models. Without this foundation, an AI solution may simply enable the business to make misguided decisions faster.

It's also critical to employ tools like GenAI with a basic idea of what 'good' looks like for the outcome you're trying to achieve. While these tools can be directed, they can't be trusted to work without supervision, or to vet the quality of the results. Having a handle on the direction and the output of your AI systems is part of a responsible technology practice, and essential to avoiding unintended consequences.

Once these parameters are in place, we encourage organizations to start testing AI with possible use cases emerging in their operations. Like all innovations, it can be difficult to understand the full potential or range of applications until the technology is firmly in play.

Signals

- **Firms discovering new applications for AI**, often in fields you wouldn't necessarily expect. Consultancy McKinsey, for example, has attracted attention with the development of a chatbot that acts as a versatile assistant and 'thought-sparring partner,' scanning the firm's vast repository of documents and interview transcripts based on employees' queries to summarize highlights and identify experts that may be able to support them in their work.
- **LLMs speaking new 'languages.'** More LLMs are being developed for specific industries or use cases that employ their own, at times highly technical, vocabulary. Examples include Bloomberg's BloombergGPT, specifically trained on financial data sets to assist with common financial services industry tasks such as named entity recognition and news classification and Google's MedPalm, which is designed to generate accurate, extensive answers to common consumer health questions and has proven knowledgeable enough to pass the US Medical License Exam.
- **AI becoming the new investment battleground.** The scale of money and resources being poured into AI by some of the world's biggest enterprises as they try to secure a competitive advantage in this fast-growing space is nothing short of staggering. Amazon's US$4 billion investment in startup Anthropic comes hot on the heels of a multibillion dollar investment by Microsoft in ChatGPT maker OpenAI earlier this year. Goldman Sachs estimates total investment in the sector could approach US$200 billion by 2025, even before AI starts to produce genuine productivity gains.
- **AI-related stocks on a tear.** The eye-popping 200%+ gain posted by NVIDIA, the leading maker of the chips that power AI systems, is just the tip of the iceberg. Shares in lesser-known firms like C3.ai, a developer of AI applications, and AI-focused exchange-traded funds (ETFs) such as BOTZ, have also surged as analysts and investors zero in on the sector.

Trends to watch

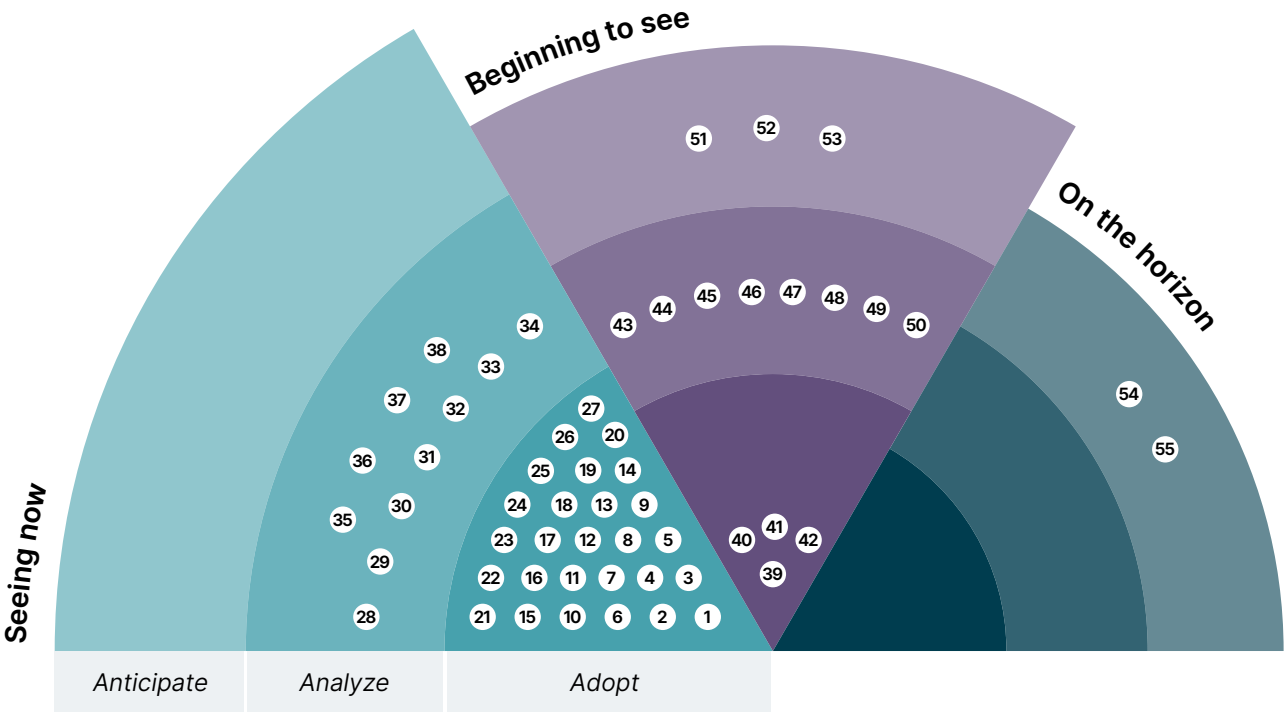

Seeing now

Adopt

1. Agent-based simulation
2. AI as a service
3. AI in security
4. AI-assisted software development
5. AI/ML on edge
6. Automated compliance
7. Code of ethics for software
8. Collaboration ecosystems
9. Data mesh
10. Edge computing
11. Ethical frameworks
12. Evolutionary architectures
13. Explainable AI
14. GenAI tools in IDEs
15. Generative AI
16. Green software engineering
17. Integrated data and AI platforms
18. MLOps
19. Natural language processing
20. Online machine learning
21. Operationalize AI
22. Platforms as products
23. Privacy first
24. RAG (Retrieval Augmented Generation)
25. Smart homes
26. Software-defined vehicles
27. Vector databases

Analyze

28. AI marketplaces
29. AI-generated media
30. AI, IoT and XR combined solutions
31. Automated workforce
32. AutoML
33. Encrypted computation
34. Federated learning
35. Personal information economy
36. Personalized healthcare
37. Robotic process automation and low code
38. Smart cities

Anticipate

—

Beginning to see

Adopt

39. AI agents
40. Causal inference for ML
41. Decentralized identity
42. Fine grained data access controls

Analyze

43. AI safety and regulation
44. Decision science
45. Digital humans
46. Easing access to Generative AI
47. Intelligent machine to machine collaboration
48. Multimodal AI
49. Production immune systems
50. Trustworthy data

Anticipate

51. Adversarial machine learning
52. Affective (emotional) computing
53. Understandable consent

On the horizon

Adopt

—

Analyze

—

Anticipate

54. AGI research
55. Quantum machine learning

The opportunities

By getting ahead of the curve on this lens, organizations can:

Smooth and accelerate human-computer interactions. Advancements in natural language processing (NLP) are opening up new ways for people to communicate with machines, including through everyday conversations. This is both broadening the scope of people who can interact with these systems and making it far easier to plan and execute tasks such as summarizing information, providing self-contained answers to inquiries or context and information to support certain roles or generating and curating content.

End the terror of the 'blank page.' Whatever the task or project, coming up with the initial idea(s) and making a start from essentially nothing is often the toughest part. AI can eliminate blank page paralysis by conducting research and providing a list of suggestions or insights that, even if far from the desired result or finished article, can serve as a jumping-off point or prod recalcitrant minds into action.

Automate tasks — not entire jobs. The need for human involvement to guide and ultimately evaluate AI output makes the wholesale outsourcing of roles to AI systems less likely than many people think. That said, there are a multitude of tasks that AI can automate or where it can augment human input, making the work more consistent and efficient. Any task that requires access to and analysis of a vast body of knowledge — such as a large number of research papers, or databases of medical or financial information — can be seen as a promising candidate for LLM assistance. The bar for what AI can do will constantly shift, but in general, as some have advised, it can be helpful to think of AI 'not as software, but as pretty good people' — that is, the equivalent of a competent research assistant or army of well-intentioned interns. There are some duties AI can be trusted with — but it certainly can't be trusted with everything.

Revolutionize software delivery. There's a misconception that in building software, GenAI's main use is as a tool to ease coding, when in reality it can touch on all aspects of the development lifecycle. Potential applications beyond coding assistance include brainstorming with AI to improve our requirements and testing scenarios; improved incident response and debugging by translating natural language into queries over logs and metrics; product and strategy ideation; and searching unstructured institutional information to provide valuable context to developers. Based on our experiences, we believe AI-assisted software delivery has the potential to drive productivity increases of up to 30%.

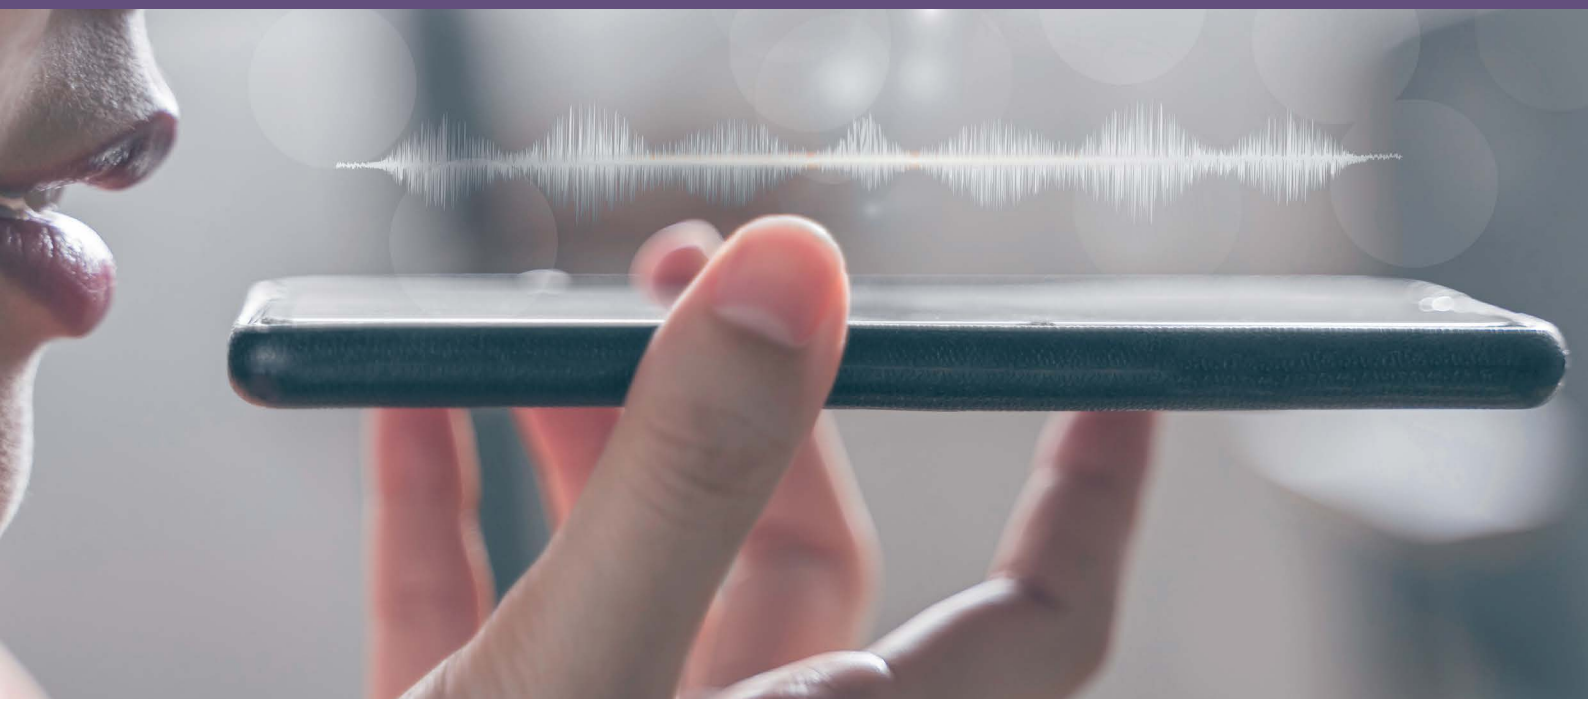

What we've done

Enhancing conversational AI with language models with Jugalbandi

We worked on a chatbot that helps users navigate the complexity of the Indian government's various welfare schemes. It's a testament to AI's ability to not only navigate, process and summarize vast amounts of information in an easily digestible format, but also to meet a much more inclusive user base on its own terms. We combined a number of existing LLMs and translation models to power conversational AI via voice, both incoming and outgoing, in multiple local languages and dialects. This provides an access point and source of information on government services to a rural user base with high illiteracy rates. This has vastly extended and simplified interactions between the Indian government and many of its citizens in remote and non-urban areas.

"As AI integration becomes more sophisticated and the implications for getting things damagingly wrong multiply, the need for effective risk management grows too."

Mike Mason
Chief AI Officer, Thoughtworks

Actionable advice

Things to do (Adopt)

- **Identify AI champions** who can help guide and teach your organization about the potential use cases for emerging solutions — but understand that AI can and will be applied in different ways in almost every part of the enterprise, which means these champions need to keep an open mind. Having people with a clear idea of what 'good' looks like can reduce risks and ensure AI initiatives focus on meaningful business results.
- **Especially in the short term, focus on how humans and AI work together.** Ensure teams understand how AI can augment, not threaten, the tasks that are core to their roles, and where their judgment will need to take over. Watch and manage the costs of services people use, which are generally not visible to an individual user but can add up quickly. Depending on the task, it can be worth sacrificing a degree of accuracy for lower costs, as models that are only slightly less accurate may be substantially cheaper to run.
- **Identify clear AI use cases that drive real value for your organization**, as well as areas where you explicitly will not use AI, either because it makes no business sense, the costs outweigh the necessary investments or the associated risks are simply too high. As the list of potential applications is massive and constantly expanding, having these decisions to orient around will ensure your AI efforts are carefully targeted and therefore more likely to bear fruit.
- **Be deliberate about which technology you're using.** 'AI' has become at times a catchall for a range of distinct technologies and, more recently, used to refer to GenAI alone. The capabilities and use cases for GenAI versus other technologies that are at times lumped under the AI umbrella — like machine learning — can be very different, and clarity is needed on what you're planning to implement and how it connects to the problems you're trying to solve.
- **Define and communicate 'guardrails' early on.** Well before they're interacting with AI on an everyday basis, teams should be aware of standards and expectations in terms of security, data sources and vetting systems or their outputs for transparency and/or bias. They should also be aware of when to give up on experiments when they are unlikely to yield the desired outcome or result in excessive risks.

Things to consider (Analyze)

- **Open-source alternatives to commercial LLMs**, which are improving and more of which are emerging every day. Freely available models can provide a solid base for chatbot and customer support applications, and be developed into specialized models that protect the organization's intellectual property.
- **AI agents.** Recent programming interfaces by companies such as OpenAI offer the ability to combine the functionality of publicly available generative AI models with specific knowledge from outside the model, such as product information.
- **New vendor offerings.** Public cloud vendors such as Amazon and Google Cloud announced a wave of new products and services for people who are creating software towards the end of 2023. In many cases these tools offer compelling features, such as AI-assisted deployment and operation of the software being created. Encourage your AI champions to regularly evaluate a variety of offerings.

Things to watch for (Anticipate)

- **Waves of regulation.** As demonstrated by the ongoing debate over an AI act in the European Union, governments are scrambling to legislate against some of AI's more negative perceived impacts and new rules about all facets of AI are likely to be coming fast from all directions. Organizations need to be proactive about establishing policies to do the right thing before they are forced to, so compliance becomes a matter of course.

Realizing value from data and AI platforms

The infrastructure supporting growth and innovation

Reliable access to credible, up-to-date data is now critical to virtually everything enterprises need to do, as underscored by the rapid rise of artificial intelligence (AI) — including generative AI (GenAI) — in business operations and decision-making. Even if organizations aren't planning on building GenAI applications in the short term, the availability of high-quality data is still a prerequisite for the delivery of other digital initiatives and services.

While data is creating unparalleled opportunities for businesses, the accompanying challenges often prove too big to ignore. Many companies remain bound by internal silos and multiple isolated data platforms that leave valuable intelligence locked up and hard to use.

Turning data stores into strategic assets requires focus on making data findable, accessible, trustworthy, interoperable and reusable, all in a secure and privacy-sensitive fashion. Data platforms provide the only viable foundation for this approach.

Turning data stores into strategic assets requires focus on making data findable, accessible, trustworthy, interoperable and reusable, all in a secure and privacy-sensitive fashion

The term "platform" means many different things, although in all cases we recommend a product thinking mindset. By integrating various data resources and ensuring they can be seamlessly accessed and applied, they provide the different building blocks now needed to form a comprehensive digital strategy.

Once in place, your data platform enables you to gather data insights, create reliable AI systems, control risk and much more. Your data platform can also be a key component in creating, managing and enforcing data governance, one of the biggest challenges faced by many organizations.

Having a robust data platform which facilitates open sharing while preserving privacy allows enterprises to participate with other organizations in thriving data ecosystems to produce greater industry, even societal, impact. This is a key trend we see expanding over the next few years that could facilitate more digital innovation and potentially create a sea change in how data is stored and exchanged — but only if a shift towards standardization takes place and enterprises learn to guard their data assets less jealously.

Signals

- **A rise of integrated data and AI platforms.** These systems present the analysis as the primary benefit, and the data just comes along. This represents a fundamental change in the way of thinking about these solutions.
- **Data ecosystems going beyond the hype to generate tangible business results.** Tech research and advisory firm Gartner sees data ecosystems moving past the peak of the hype cycle and firmly entering the mainstream within the next decade. Our experience on the ground points to the same trend, with growing demand for organizations to share and pool data resources — enterprises are today proving more willing to take the plunge.
- **Data 'clean rooms' gaining popularity.** A number of vendors including Infosum, AWS, Google and Snowflake have developed data 'clean room' offerings designed specifically to enable the secure intra- and inter-organizational sharing of privacy-compliant, anonymized data efficiently and at scale.
- **The creation of open data sharing standards and infrastructure gathering momentum.** As exhibited by initiatives like the Open Data Standard for the Apparel Sector and OSDU™ Forum, some organizations are attempting to circumvent the data interoperability challenges perpetuated by incumbent vendors by championing open protocols for the secure exchange of valuable data on supply chains and other industry-critical functions.
- **Mechanisms for privacy-aware sharing of user data.** Data privacy is increasing its mindshare among users. We have developed a solution called Anonymesh to help address this challenge. Organizations are also building personal data stores such as Solid Pods that organize data storage around users, instead of the organizations that collect them.
- **Governments setting the open data agenda.** Countries worldwide are embracing open data initiatives that encourage free access and the use of government-collected information to improve public services and create economic opportunities. The UK government, for example, is building an Integrated Data Service (IDS) to facilitate cross-departmental data exchange. Meanwhile, Singapore's open data initiative enables developers to use real-time datasets from government agencies to develop their own applications.

Trends to watch

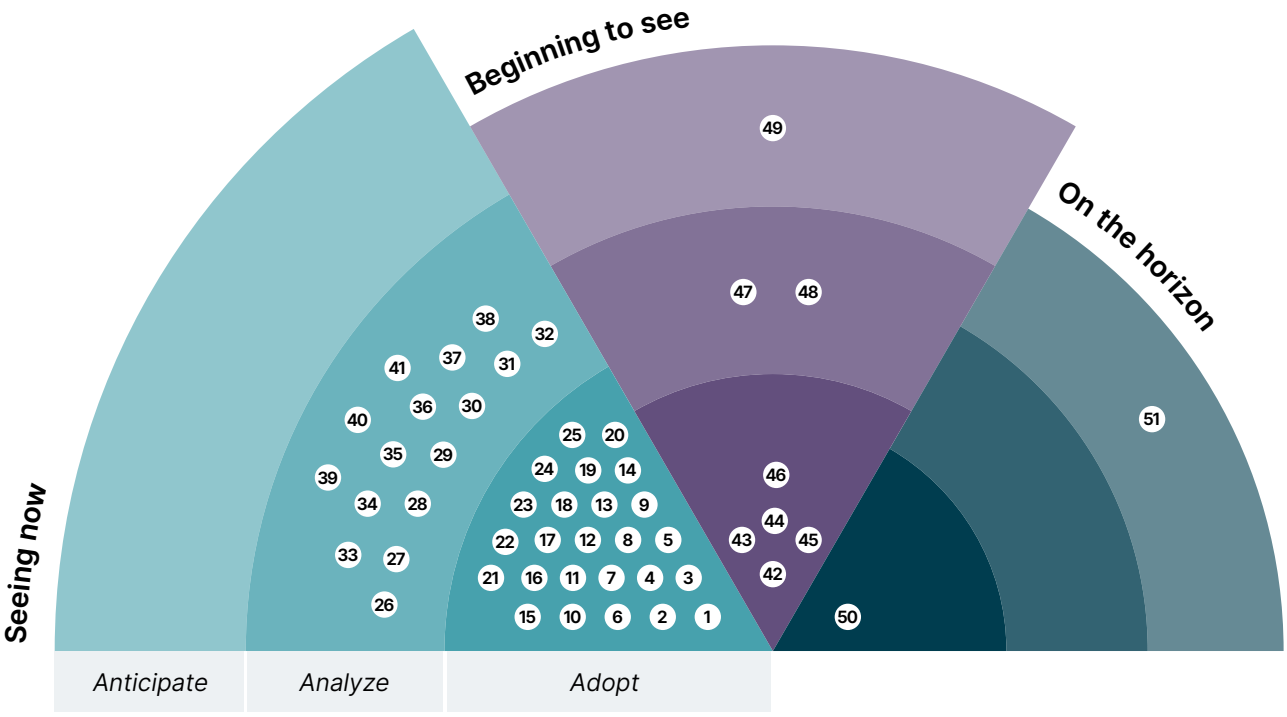

Seeing now

Adopt

1. AI as a service
2. AI/ML on edge
3. Automated compliance
4. Code of ethics for software
5. Collaboration ecosystems
6. Data mesh
7. Digital ecosystems
8. Digital twin
9. Edge computing
10. Ethical frameworks
11. Explainable AI
12. FinOps
13. Green cloud
14. Integrated data and AI platforms
15. Knowledge graphs
16. MLOps
17. Online machine learning
18. Platforms as products
19. Privacy first
20. Privacy-enhancing technologies (PETs)
21. RAG (Retrieval Augmented Generation)
22. Secure software delivery
23. Smart homes
24. Smart systems and ecosystems
25. Vector databases

Analyze

26. AI, IoT and XR combined solutions
27. AutoML
28. Autonomous robots
29. Autonomous vehicles
30. Data clean room
31. Data marketplaces
32. Decentralized data platforms
33. Differential privacy
34. Encrypted computation
35. Federated learning
36. Increased regulation
37. Personal information economy
38. Privacy-aware communication
39. Privacy-respecting computation
40. Smart cities
41. Zero knowledge proofs

Anticipate

—

Beginning to see

Adopt

42. Causal inference for ML
43. Data contract
44. Decentralized identity
45. Fine grained data access controls
46. Re-decentralization

Analyze

47. Intelligent machine to machine collaboration
48. Trustworthy data

Anticipate

49. Responsible tech facilitation

On the horizon

Adopt

50. Data product specification

Analyze

—

Anticipate

51. Decentralized personal data stores

The opportunities

By getting ahead of the curve on this lens, organizations can:

Enable AI and GenAI initiatives. High quality data is a fundamental requirement for any artificial intelligence initiative. Forbes called data quality "the real bottleneck in AI adoption".

Improve compliance posture and reduce risk. With automated embedded governance policies created and enforced via the underlying platform, you reduce the gap between written polices and what actually gets implemented on the ground.

Achieve significant savings by eliminating redundancy. Enhanced data sharing and creating a platform that provides a single point of availability allows organizations to retire tech infrastructure that's duplicated across different parts of the organization, substantially reducing the costs of technology infrastructure and upkeep.

Gain a competitive advantage with improved insights. Integrating quality data across the enterprise can highlight previously unnoticed areas of inefficiency or friction, as well as provide a more holistic view of complex processes like the supply chain or customer journey. The resulting context and discoveries can help the organization understand their workflows and customers better, granting the enterprise an edge over competitors whose data resources remain more piecemeal and isolated.

Develop new sources of value. As data interoperability capabilities mature and open standards gain more widespread acceptance, more opportunities to capitalize on data assets will arise in the form of data marketplaces and networks. While some large organizations and industries have begun creating their own data ecosystems — such as those emerging around open banking and data sharing in insurance — there is still room to expand these to other sectors and the wider community.

Ramp up time-to-market. The availability of high quality data, especially as a basis for GenAI, is poised to accelerate and enhance many painstaking aspects of the product development process. One example is how the ability to sift through and summarize vast amounts of information and create synthetic customer data is helping organizations substantially reduce the lead time needed for market research. We believe large enterprises will have an edge over startups in pursuing these opportunities if their vast data reserves are used effectively.

What we've done

Helping ITV better leverage data with a data mesh platform

In response to the paradigm shift of digital streaming, ITV set a new vision: to become a digitally-led media and entertainment company that creates and delivers standout content to audiences wherever, whenever and however they choose. An expert team of Thoughtworkers began co-developing a cloud-based data mesh on AWS and Databricks, a process that would enable ITV to bring its new data strategy to life and embed agile ways of working across its diverse business units.

ITV's data mesh platform enables teams to quickly onboard their data and make it discoverable and easily accessible across the business. The time taken to provision data products using the platform has gone from three weeks to just a few hours — driving the adoption and expansion of the mesh across ITV's operations.

"Data doesn't carry any intrinsic value on its own. Its value relies on you having a purpose — and a process — for it."

Emily Gorcenski
Head of Data & AI, Thoughtworks Europe

Actionable advice

Things to do (Adopt)

- **Implement privacy-enhancing technologies (PET).** These technologies provide increased privacy or secrecy for the persons whose data is processed, stored and/or collected by software and systems. They are often used as a part of this processing and modify the normal ways of handling (and often hoarding) raw or plaintext data directly from users and internal participants, such as employees. By increasing the privacy offered, you are both reducing the owned risk and providing users with better choices about how they'd like their data to be treated.
- **Enhance data governance and privacy policies.** If data is not sufficiently protected and governed internally with clear principles around issues such as privacy and consent, it becomes too risky to take any steps that expose it to the outside world. Before considering wider data sharing and collaboration, enterprises need to clearly define the scope of access and influence over data that various roles have, and embed compliance with policies-as-code into their data platforms.
- **Streamline data processes and the path to production.** Practices like DataOps and MLOps offer techniques to speed up key aspects of the production cycle and improve developer experience, with shorter feedback loops and guardrails that ensure risks are still mitigated.
- **Embrace data mesh to deliver insights at scale.** Experimenting with data mesh architecture can provide the integration and accessibility needed for various teams to make the most of the data in their domains. This will improve visibility over processes and give teams the ability to rapidly direct development to serve business needs, helping future-proof the enterprise.

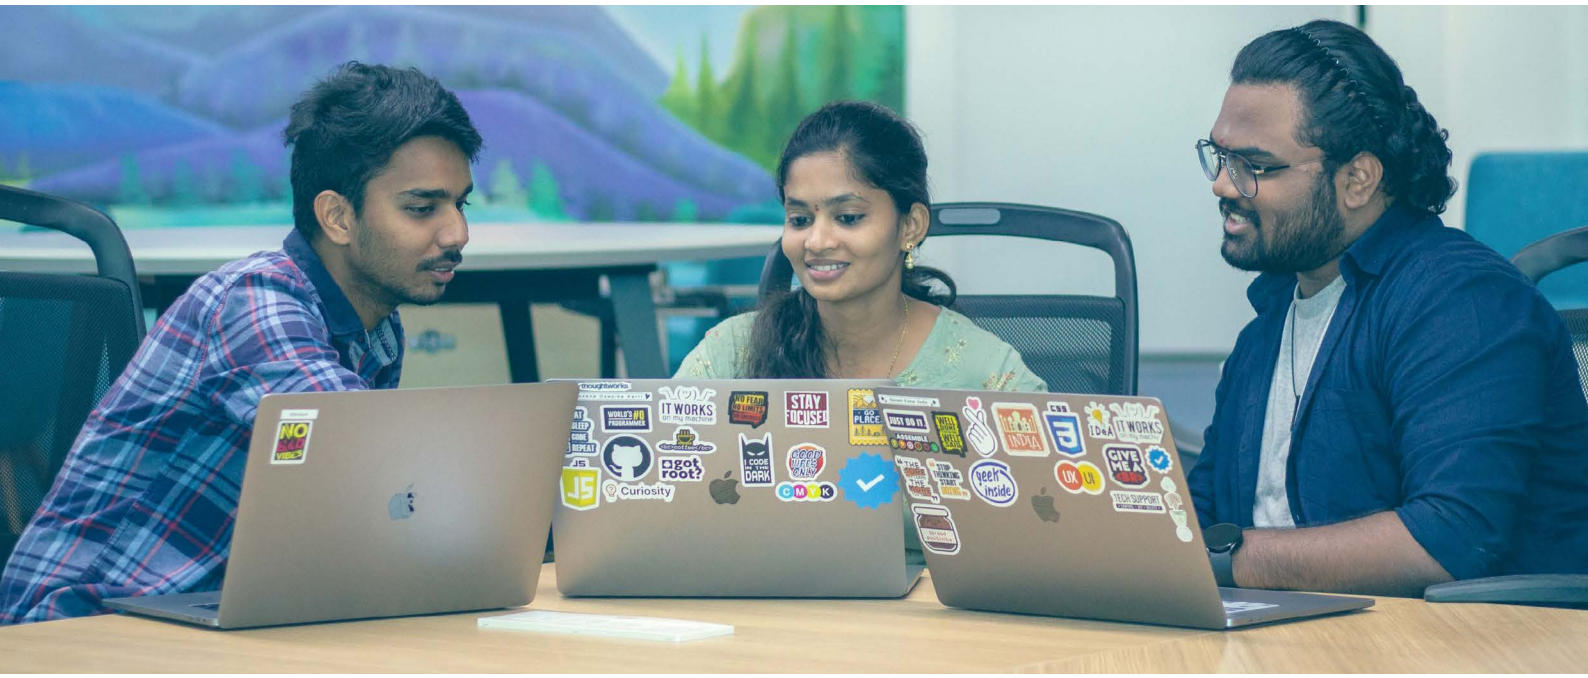

Things to consider (Analyze)

- **Storing data differently.** Emerging trends like data clean rooms and differential privacy, which preserves the anonymity of individual aspects of a data set by introducing 'noise' around it, can provide a stronger basis for the enterprise to house data in a trustworthy and compliant way while remaining able to put it to use.
- **Participating in data marketplaces.** As more examples of open, pooled data marketplaces emerge, such as the version being advanced by the UK government, businesses should consider their appetite for participating in these initiatives and what the potential of participating in them might be. It's important to examine questions such as: where might ecosystems enable the enterprise to create more value? And what, if any, capacity is there for the business to monetize data while remaining sensitive to security and customer privacy?
- **Utilizing data product specification and data contracts.** Open specifications which aim to set out and standardize how data is shared between or consumed by various parties are gaining traction, and they may need to be integrated into the organization's data platform and wider strategy.

Things to watch for (Anticipate)

- **The dovetailing of data and responsible technology practice.** Responsible tech principles provide an increasingly valuable roadmap for enterprises keen to extend and make optimal use of their data resources in a consistently ethical way.
- **Decentralized personal data marketplaces** that give consumers more sovereignty over their personal data are contributing to the development of a personal information economy. This trend will have significant implications for the way companies store, analyze and use information about their customers, and the subsequent development of enterprise ecosystems.

Evolving interactions

New opportunities to engage and interact

The range of methods for human-computer interaction is expanding, and interactions themselves are growing more immersive and seamless. This is creating new opportunities for organizations to reimagine how they engage with, learn from and delight their customers, employees and other stakeholders.

Many of the trends in this area are easy to overlook as they don't, at least on the surface, seem revolutionary. In fact, the accuracy and usefulness of these systems has and will continue to grow immensely in the next couple years. For example, users have been able to speak commands into their mobile devices for years, but, until recently, interaction had to be phrased in a way the device would understand.

Beyond voice and text, we expect continued advancement in extended reality (XR) technologies that allow users to interact in virtual worlds, though perhaps not at the pace predicted by the early champions of the space

Products and services such as OpenAI's ChatGPT, Google's Bard and Microsoft's Copilot have leveraged advances in Generative AI to lead this charge, lowering the bar for individuals to interact with systems or directly with AI models. Tech giants such as Apple are also undertaking major overhauls of their voice-to-text systems, using better predictive AI and context-aware models to render technology touch-free. Beyond voice and text, we expect continued advancement in extended reality (XR) technologies that allow users to interact in virtual worlds, though perhaps not at the pace predicted by early champions of the space.

There's significant potential for these evolving modes of interaction to make technology experiences more inclusive, as demonstrated by systems like Jugalbandi which enables illiterate people to access information about government programs by speaking to their mobile device in their native language. Yet they also present a broader set of risks and accessibility issues. While accessibility approaches for more traditional interactions like mobile are relatively well understood, this is not the case for newer interactions like XR. Making voice- or GenAI-based interactions more available and sophisticated widens the scope for misuse and unintended consequences.

For the near term, there is no shortage of promising use cases for XR in emerging areas like training and data visualization that all enterprises should be ready to explore. Consumer adoption is likely to remain primarily limited to areas like collaboration, gaming and entertainment. So far, the advancements in consumer devices haven't been sufficient to expand outside these areas, but product development by Apple, Meta and others is expected to continue.

Signals

- **Tech giants unveiling new XR-enabled devices** such as Apple's Vision Pro 'spatial computer,' and Meta's Ray Ban smart glasses. These are designed to be more comfortable and less intrusive than earlier examples of such hardware. In Apple's case it incorporates sophisticated gesture recognition to reduce the need for physical input devices.
- **The development of GenAI-enabled apps** that enhance accessibility for people previously facing physical and/or linguistic barriers to working with technology systems. Jugalbandi, for example, enables interactions in multiple Indic languages simply through the user's voice. It does this using a combination of ChatGPT, language translation models and a fixed data set from which to answer questions.
- **Higher education institutions pushing the frontier of data visualization.** Researchers at institutions such as Monash University are experimenting with using XR to bring representations of data beyond screens into more immersive hybrid environments.
- **The release of enhanced libraries for gesture recognition** designed to help developers connect physical gestures to application features, such as those developed by Google and Apple.
- **Continued advancements in natural language processing**, including the development of pre-trained models to address specific tasks like sentiment analysis, and a growing focus on multilingual capabilities. These advancements are being leveraged by forward thinking organizations to improve how users find and use information. For example, Zalando has created an interactive assistant that goes far beyond what's possible with parametric search.

Trends to watch

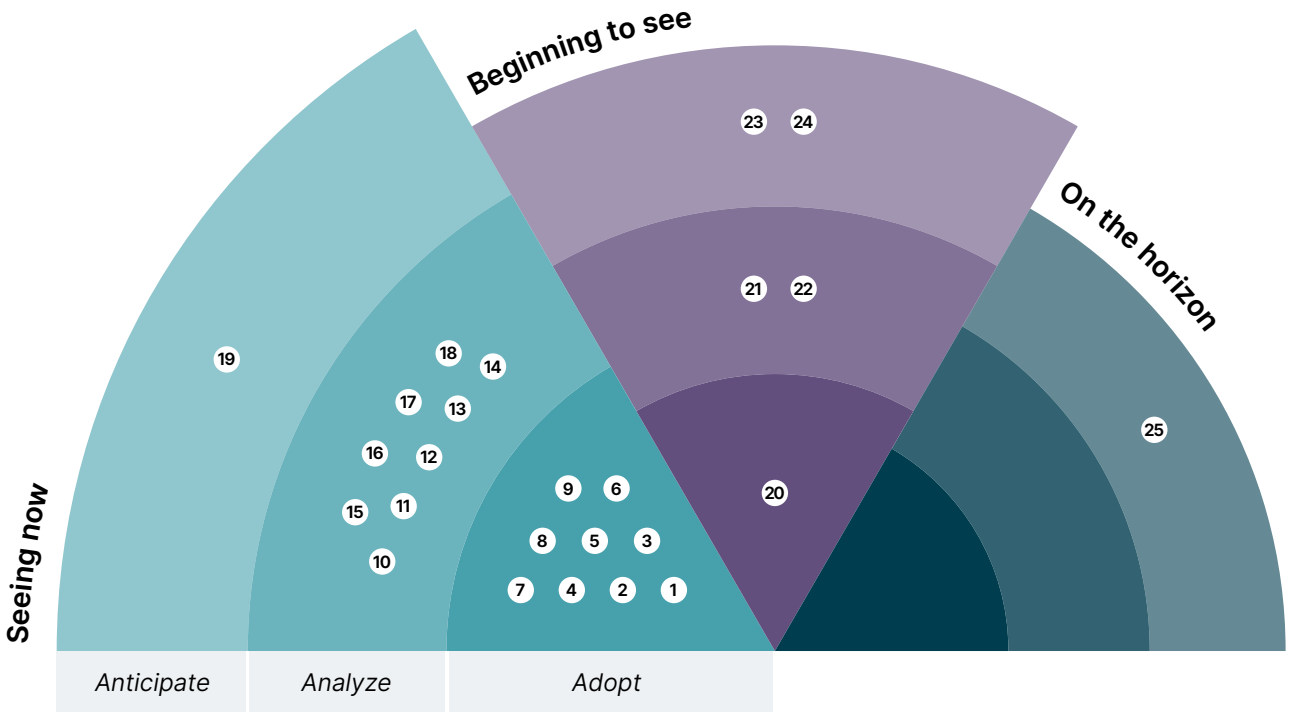

Seeing now

Adopt

1. AI-assisted software development
2. Developer experience platforms
3. Enterprise XR
4. Natural language processing
5. Operationalize AI
6. Privacy first
7. RAG (Retrieval Augmented Generation)
8. Smart homes
9. Ubiquitous connectivity

Analyze

10. Augmented reality
11. Automated workforce
12. Consumer XR
13. Gesture recognition
14. Personalized healthcare
15. Satellite networks
16. Smart cities
17. Spatial audio
18. Touchless interactions

Anticipate

19. Addictive tech

Beginning to see

Adopt

20. Industrial XR

Analyze

21. Digital humans
22. Retina resolution XR

Anticipate

23. Understandable consent
24. XR-enabled hybrid working

On the horizon

Adopt

—

Analyze

—

Anticipate

25. AGI research

The opportunities

By getting ahead of the curve on this lens, organizations can:

Lower costs by making interactions more efficient. The ability to communicate with a system with the flick of a finger or through voice commands in natural speech, rather than the more laborious process of typing, has the potential to massively accelerate productivity in the workplace and other environments.

Boost satisfaction by reducing friction and enhancing availability for customers. Voice-based platforms promise to make it easier for customers, especially those with accessibility challenges, to retrieve product information or get answers to queries in a smooth, hassle-free way. Similarly, advances in natural language processing opens up possibilities such as using chatbots to provide a base level of service and support outside working hours.

Derive better insights from richer interactions. As AR and VR make it possible to bring data to life outside the screen or printed page, the foundations are being laid for what's known as immersive analytics. Immersive environments can help users experience data aurally and tactically as well as visually, enhancing understanding, analysis and, eventually, decision-making.

Test out scenarios to improve responses. XR-enabled simulations can be used by enterprises to run teams through mission-critical situations that could test the business, giving them an accurate sense of their capacity to react and identify areas for improvement.

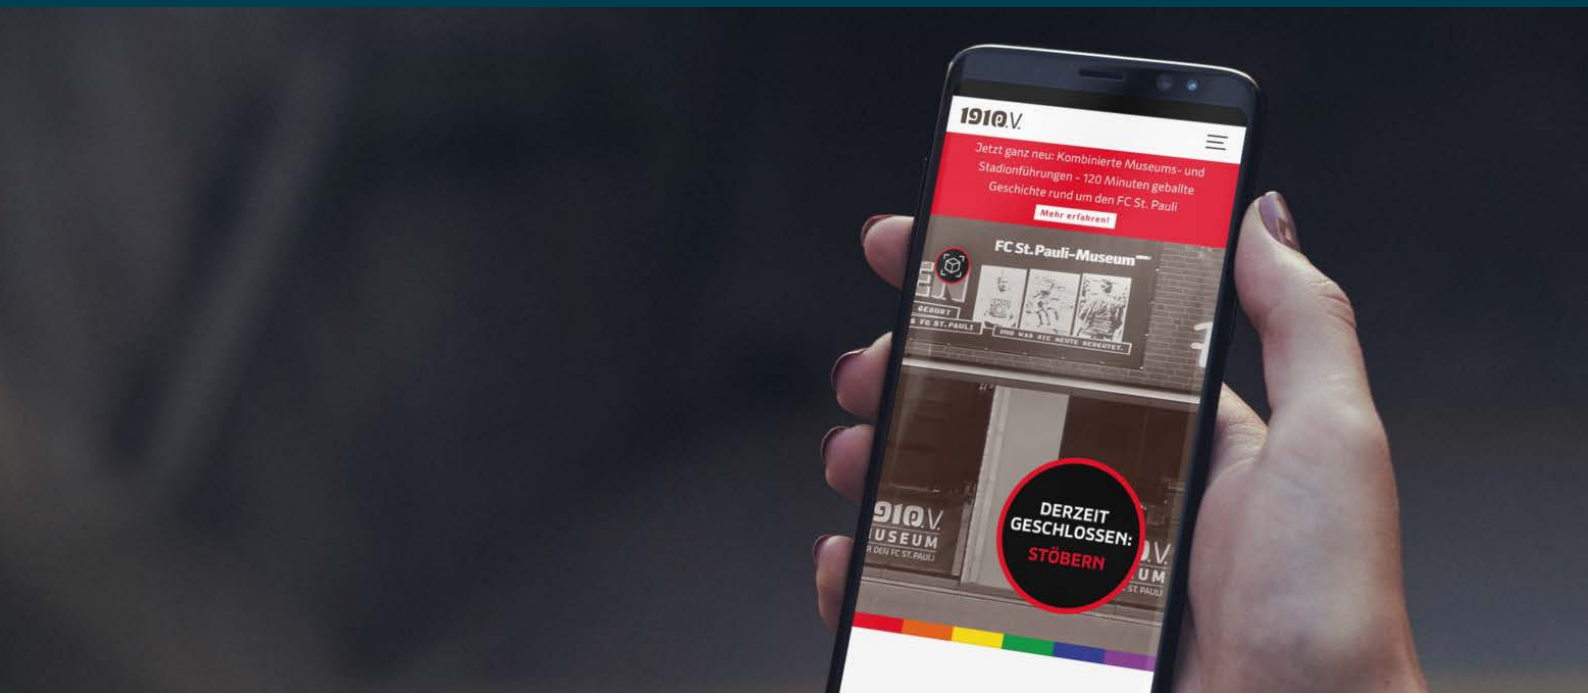

What we've done

An immersive fan experience for the FC St. Pauli Museum

The Hamburg-based football club museum's ambition is to create a dynamic and immersive experience that grows from fan engagements and promotes topics such as diversity, sporting events against racism, and wider conversations around social impact. Thoughtworks worked closely with FC St. Pauli Museum on a six-week project to design a new interactive fan experience as part of Tech Lab — a Thoughtworks initiative that takes innovative ideas and uses cutting-edge technology to bring them to life.

"You often see people's eyes rolling at terms like metaverse or XR, but at the same time, people are using these immersive technologies in their daily lives without even realizing it: just look at how readily people accept the blurring or changing their surroundings on video calls."

Kuldeep Singh
Principal Consultant, Thoughtworks

Actionable advice

Things to do (Adopt)

- **Actively investigate processes or areas that can be improved or replaced by advances in AI and interaction technologies, such as:**
  - **Customer interactions,** more aspects of which can be supported by chatbots as they grow more sophisticated and attuned to industry-specific use cases.
  - **Researching and uncovering market trends.** Business analysts at Thoughtworks have been able to use ChatGPT and other tools such as Boba for ideation and scenario creation.

Things to consider (Analyze)

- **Monitor the XR/AR space for possible use cases.** The benefits for things like training and crisis management, where realistic physical interactions are particularly important, could be substantial. In addition to enterprise contexts, organizations derive value from industrial use cases ranging across design, manufacturing and maintenance applications. The cutting-edge work taking place on data visualization is also likely to prove important for how we consume and manipulate data in future.
- **Understand the diversity of AI tools** and how this can serve the needs of people who may want to interact with AI in different ways. Some developers, for example, are comfortable with tools like GitHub Copilot where the actual interaction is largely the same as it has been for some time with autocomplete functions. Others would rather have a 'conversation' with a ChatGPT-like solution and use the resulting learnings in their regular internal development environment.
- **Learn, and go, where your customers are.** Consider your customer base and how likely your average users are to adopt any new platforms for interaction you offer. For example, retailers could focus on augmented reality (AR) interactions which enable richer interactions for consumers in physical locations.

Things to watch for (Anticipate)

- **Consider 'outside the box' use cases.** There's no question the expanding nature of interactions will result in potential applications for every business, but these applications may not always be obvious. This means it's important to think through business processes and examine where the ability to present data or engage with customers in an entirely different way might improve the entire experience.

Accelerating physical-digital convergence

Closing the gap between the real and the virtual

While automation has been a presence in industries like manufacturing and agriculture for decades, never has its reach and impact extended so far into the physical world — and the trend is only set to accelerate.

The previous generation of automation was contained within embedded systems that were often subject to physical constraints; for example, robots confined to a single assembly line, or those that struggled with pathfinding when confronted by a wall. But newer sensors and tools are capable of not just mapping and penetrating physical space over a far wider area, but of also recreating that space digitally with astonishing fidelity.

Newer sensors and tools are capable of not just mapping and penetrating physical space over a far wider area, but of also recreating that space digitally with astonishing fidelity

Today, this lens lends itself to a variety of applications that are likely to be most relevant to organizations with a clear connection to the physical environment, but we expect it to impact any organization which is not a pure digital play. New agriculture systems capable of measuring the soil and applying water and fertilizer in optimal proportions are improving crop yields, demonstrating the extent and precision to which the digital can now alter, and even enhance, the physical world. Digital twins, or detailed digital representations of physical objects in their real-world context, are being used in sectors like aerospace to monitor and evaluate the performance of equipment, flagging potential problems before they appear in reality.

This convergence of environments requires more holistic thinking. High-quality data needs to be available and shared, because systems won't be able to make intelligent decisions without it. In the past, engineers working on embedded systems would most likely sit in one building while those tasked with building an engine or vehicle worked in another, but now functions should come together to build and run products that span both worlds.

Signals

- **The expansion of autonomous vehicles.** While not without controversy or setbacks, driverless vehicles continue to take to the roads in greater numbers and in more locations. Robotaxi pioneers Waymo and Cruise are reportedly mulling expansions to Los Angeles, Texas and even Japan. They're also pushing the development of new sensor technologies in the process.
- **Smart farming that drives results.** Startups like Israel's BloomX are blending digital and physical systems and blanketing farms in technology to mimic and improve on natural processes like pollination. These innovations have the potential to improve output and mitigate some of the challenges climate change poses to the agricultural sector.
- **The emergence of new platforms in manufacturing.** Research from analysts such as Gartner suggests the density and variety of robots in the typical warehousing or manufacturing operation is set to surge, with nearly all such firms planning to expand their robot workforces. As these 'fleets' are made up of machines from various vendors, a category of software is emerging that is capable of integrating and communicating with all of them to coordinate their work. This is what Gartner calls 'multi-agent orchestration platforms.'
- **Even more proactive maintenance.** Developments in artificial intelligence are elevating the practice of predictive maintenance, with companies like Shell applying AI to analyze historical data and real-time sensor readings to assemble a more granular picture of the health and performance of their assets. This is further enhancing their ability to identify points of failure before failures actually occur.

Trends to watch

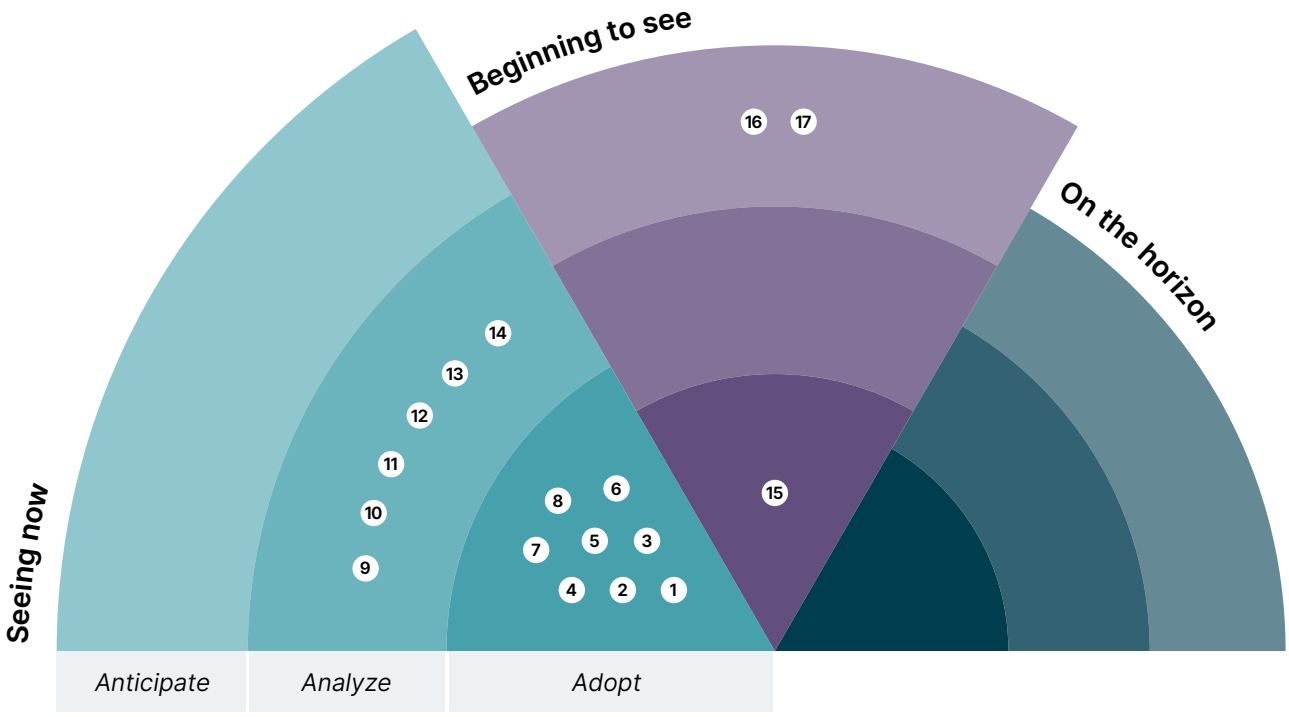

Seeing now

Adopt

1. DevSecOps
2. Digital carbon management
3. Digital twin
4. Distributed energy resources
5. Operationalize AI
6. Privacy first
7. Smart homes
8. Software-defined vehicles

Analyze

9. AI, IoT and XR combined solutions
10. Augmented reality
11. Gesture recognition
12. Personalized healthcare
13. Satellite networks
14. Touchless interactions

Anticipate

—

Beginning to see

Adopt

15. Industrial XR

Analyze

—

Anticipate

16. Affective (emotional) computing
17. Brain computer interfaces

On the horizon

Adopt

—

Analyze

—

Anticipate

—

The opportunities

By getting ahead of the curve on this lens, organizations can:

Conserve resources and reduce risks. All indications suggest that moving towards more real-time monitoring of assets and using data analysis or simulations to identify potential weak links in the production process can have a near-immediate impact. One study by PwC of companies in Europe found the average firm implementing predictive maintenance registered a 9% bump in uptime and a 12% reduction in costs, as well as a double-digit decrease in risks like health and safety.

Smooth the path to better output. Beyond boosting yields by making the process more efficient and less error-prone, as seen in the case of smart farming, the physical-digital convergence is optimizing the production process by vastly reducing the costs and complexity of prototyping and testing.

Enhance, and respond to, customer experience. While full-scale automation may provoke more skepticism, studies show consumers are responding positively to targeted technology-based solutions that reduce challenges in the physical world, such as blind spot warning and lane assist systems in vehicles, or smart homes that promote energy efficiency. Data from the use of such systems can also provide insights into customer behavior and engagement that can be used to refine products further.

Contribute to sustainability goals. Hybrid physical-digital environments promise to help firms reduce their environmental impact in myriad ways, whether by optimizing the use of resources like land, water and energy; reducing the number of times parts or machines need to be trashed and replaced; or minimizing the need to run real-world tests on physical objects or prototypes.

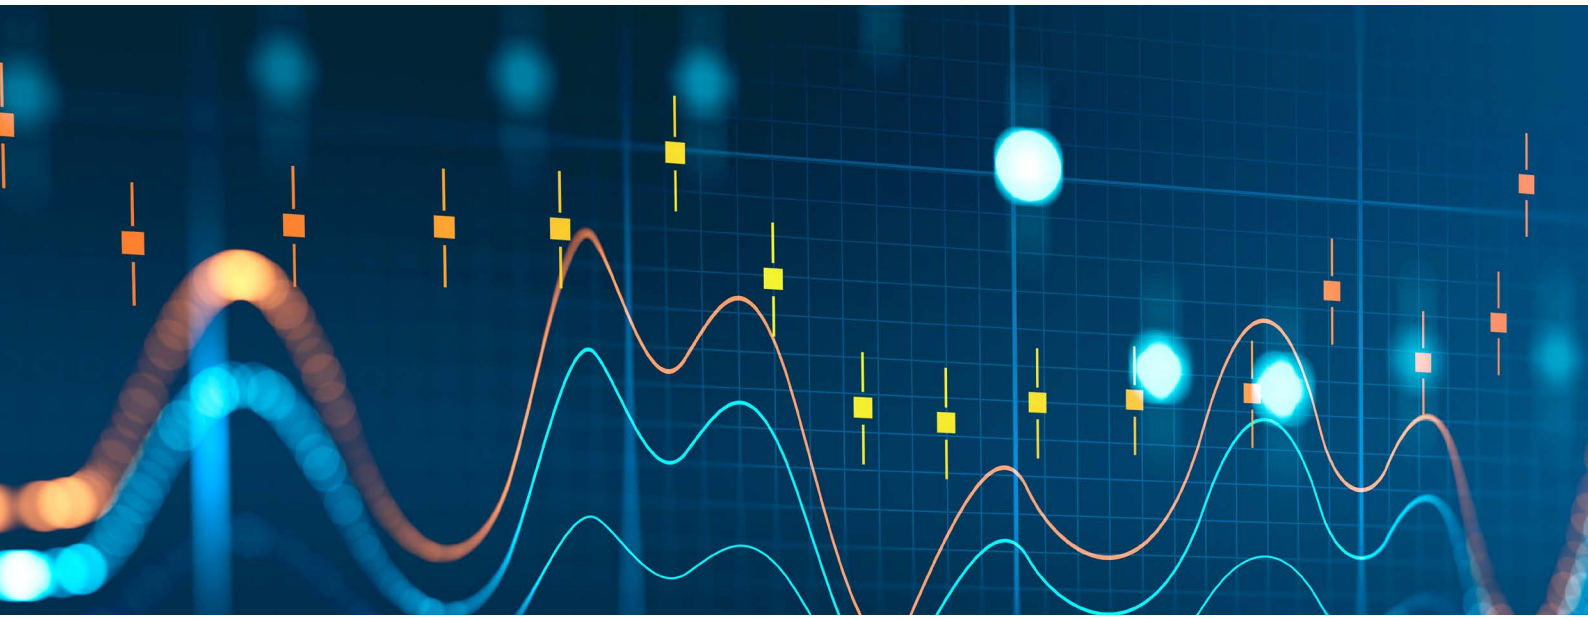

What we've done

Making autonomous vehicles a reality with MOIA

Though autonomous vehicles face various physical, regulatory and psychological barriers to widespread adoption, progress is being pushed by innovative firms like this Volkswagen-affiliated ride-pooling provider, which enlisted our help to test out scenarios to make autonomous vehicles part of their fleet. After developing a number of simulations that extended MOIA's capabilities to deploy and manage these vehicles quickly, safely and seamlessly, Thoughtworks is now supporting their rollout to real-world settings, using data and devices to measure how they respond to complex road conditions and traffic challenges.

"We're past the inflection point in the adoption curve of connected devices: we've gone through some of the early crazy IoT ideas, the pushback and reached that point where we're seeing genuinely useful innovations."

Michael Fait
Head of Software-Defined Vehicles, Thoughtworks Germany

Actionable advice

Things to do (Adopt)

- **Understand that physical-digital convergence applies to you,** unless your enterprise is a pure digital play (that is, produces no physical output). There will be use cases for your organization; the question is where the most relevant ones lie and which should be implemented first.
- **Consider the evolving digital sphere when thinking about your physical footprint.** Many of the people working on physical tasks that could formerly be performed largely in isolation and in a self-contained way — for example, planting explosives in construction or drilling in mining — will need to be connected to and interact with teams creating software and designing systems that manage the increasingly essential digital aspects of these processes.

Things to consider (Analyze)

- **Look for opportunities to improve physical processes** by taking advantage of how technologies are developing. Sensors and actuators, in particular, are far more precise and powerful than they were even a few years ago. Companies are focused not just on making them smaller, so they can be embedded in even more devices, but ensuring they can operate independently and deliver more granular data.

Things to watch for (Anticipate)

- **Monitor compliance and regulations** around fields like automation and the internet of things (IoT), which are likely to increase as digitally-enabled devices grow more sophisticated, intelligent and potentially intrusive.
- **Pay attention to the moves and initiatives of tech firms, governments and peers in the space.** As there is only one physical reality, organizations will effectively be forced to adopt whatever technologies or protocols emerge as standards are established in areas like data transfer and sensor communications. Adopting and sticking to the wrong standards or technologies could see the enterprise effectively cut off.

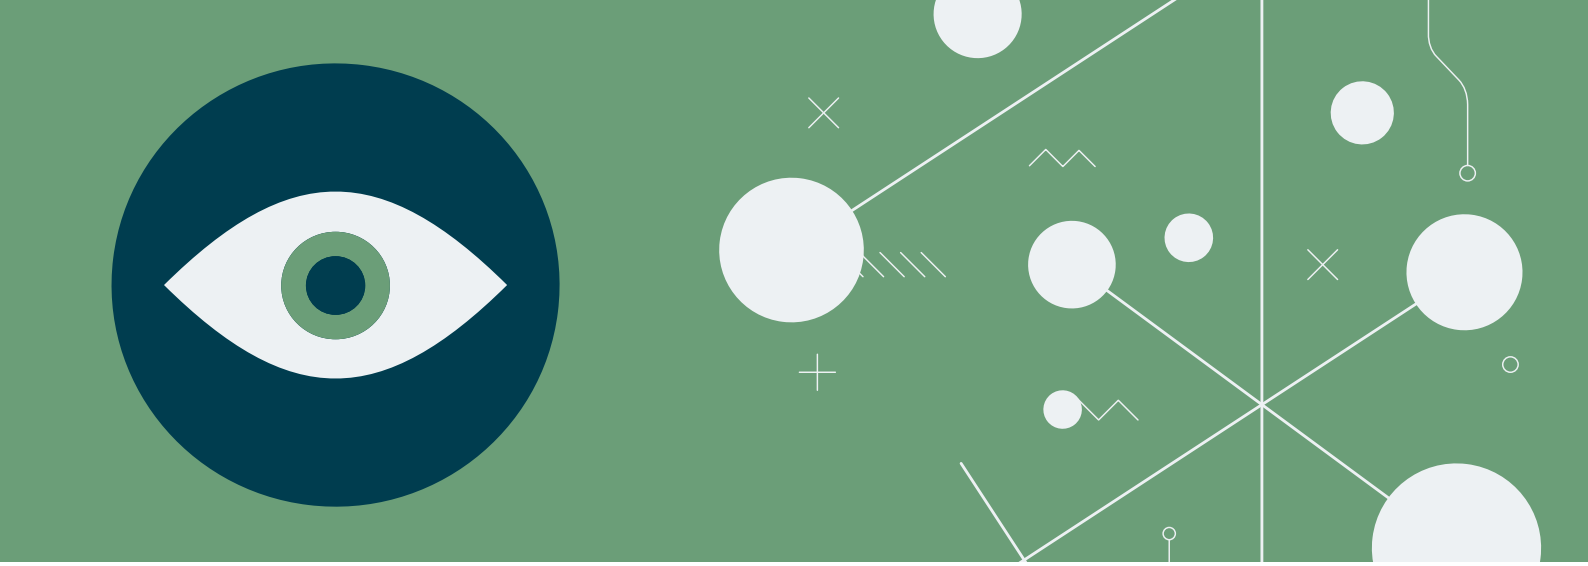

Responsible tech: a critical consideration

Building ethical guardrails

As technology has become more deeply interwoven in various parts of everyday life, the possibility of harm — both intentional and inadvertent — becomes more acute. The ongoing debate about the far-reaching implications of generative AI (GenAI) is just one example of how focus on the potential fallout from technology solutions, whether in the form of misinformation, excessive carbon emissions or the exclusion of certain groups, is increasing. According to our research on what consumers want from GenAI, 93% of those surveyed say businesses which fail to incorporate responsible and ethical thinking risk detrimental impacts.

Organizations must be prepared for their technology practices to come under more scrutiny and think through the ethical ramifications of their technology choices — not just for end-users, but for society as a whole.

Responsible technology ensures consideration for all stakeholders, as well as guardrails around privacy, security and sustainability, are firmly embedded in the organization's technology approach.

As leaders we must recognize that we often struggle to accurately predict the consequences of our technology choices

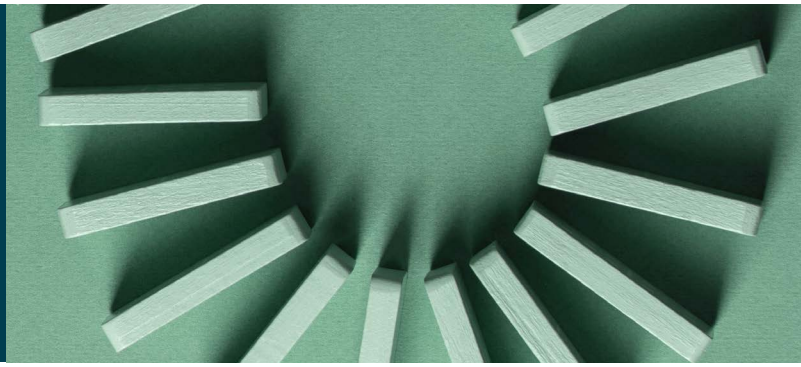

As leaders we must recognize that we often struggle to accurately predict the consequences of our technology choices. More often than not the negative effects of technology are unintentional — but that makes them no less harmful. Practicing responsible technology is a matter of broadening our perspectives and taking advantage of emerging tools and techniques that can support enterprises on their responsible tech journey, from secure software to privacy-first design.

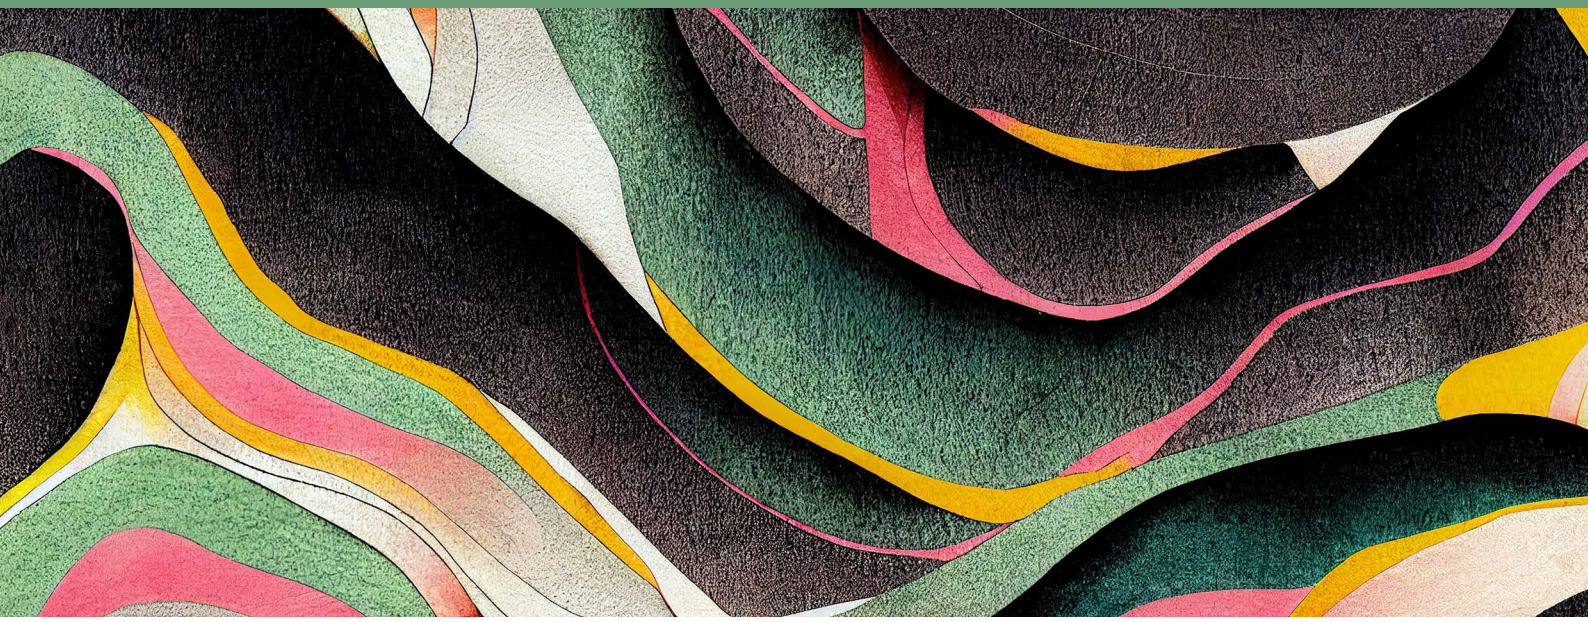

Signals

- **New resources to define and provide education on responsible technology practices**, such as the Responsible technology playbook developed jointly by Thoughtworks and the United Nations, as well as the Social and Ethical Responsibilities of Computing curriculum developed by MIT, available via the institution's OpenCourseWare platform.
- **Regulatory and policy initiatives designed to mandate more considered approaches to technology**, such as the US Executive Order on the Safe, Secure, and Trustworthy Development and US of Artificial Intelligence, India's Digital Personal Data Protection Act 2023, and the EU's proposal for harmonized regulation on AI and connected impact assessment.
- **The emergence of alliances focused on developing and promulgating sustainable technology practices.** The Green Software Foundation, which has developed training and code for less carbon-intensive software and methodologies to calculate the emissions associated with technology, is one prominent example.
- **The birth of investment funds targeting responsible technology companies and solutions**, such as Mozilla Ventures, which is channeling $35 million to early-stage startups working in privacy, decentralizing digital power and ethical AI.
- **An uptick in firms promoting responsible technology principles and credentials**, including giants like IBM, PwC and Salesforce.

Trends to watch

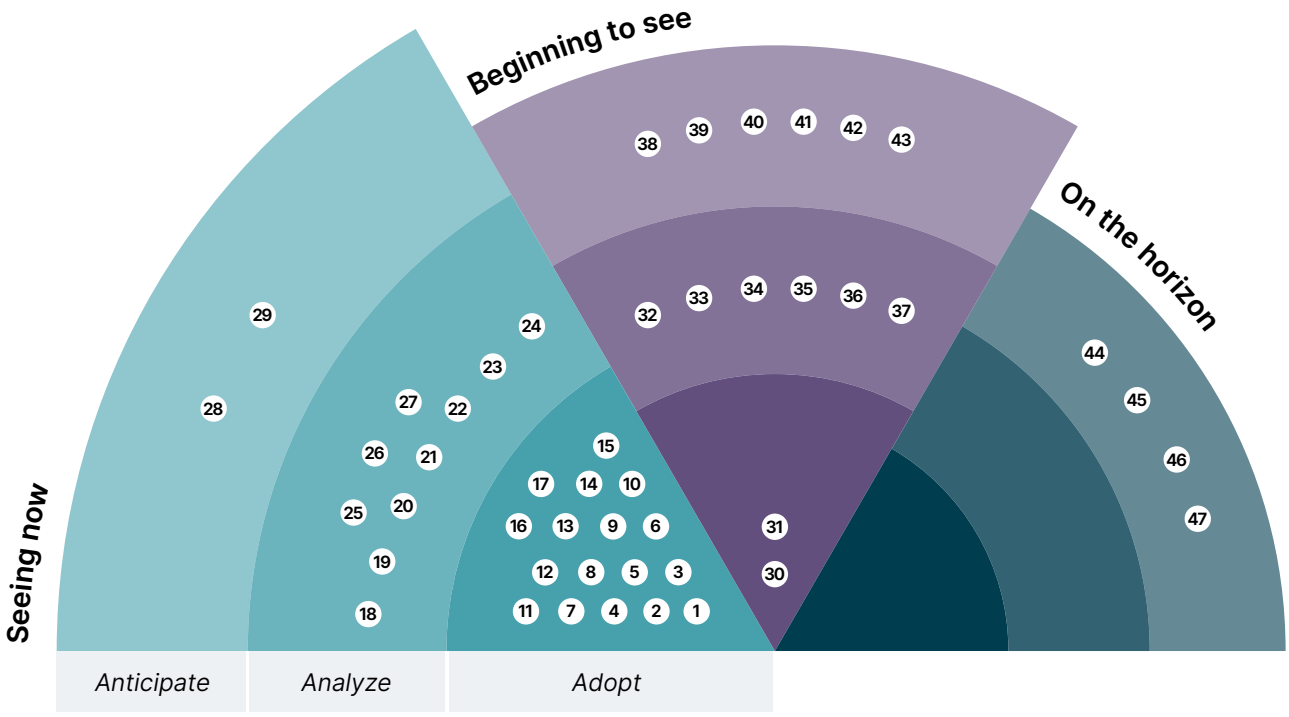

Seeing now

Adopt

1. AI in security
2. AI-assisted software development
3. Automated compliance
4. Code of ethics for software
5. Decentralized security
6. DevSecOps
7. Digital carbon management
8. Distributed energy resources
9. Edge computing
10. Ethical frameworks
11. Green cloud
12. Green software engineering
13. Privacy first
14. RAG (Retrieval Augmented Generation)
15. Secure software delivery
16. Smart homes
17. Vector databases

Analyze

18. AI-generated media
19. Alternative currencies
20. Automated workforce
21. Data marketplaces
22. Differential privacy
23. Encrypted computation
24. Smart cities
25. Smart energy management systems
26. Technology and sovereign power
27. Zero knowledge proofs

Anticipate

28. Addictive tech
29. International law for crypto assets

Beginning to see

Adopt

30. Decentralized identity
31. Fine grained data access controls

Analyze

32. AI safety and regulation
33. Decision science
34. Digital humans
35. Green UX
36. Production immune systems
37. Trustworthy data

Anticipate

38. Adversarial machine learning
39. Affective (emotional) computing
40. Quantum computing
41. Responsible tech facilitation
42. Technology for circular economy
43. Understandable consent

On the horizon

Adopt

—

Analyze

—

Anticipate

44. AGI research
45. Decentralized personal data stores
46. Next-generation cryptography
47. Quantum machine learning

The opportunities

By getting ahead of the curve on this lens, organizations can:

Prevent reputational damage among customers, talent and investors when technology-driven ethical lapses enter the public sphere. Beyond mitigating harm, an effective responsible technology practice can pay dividends in terms of customer and talent attraction and retention. One recent survey of millennial and Gen Z workers, for example, found they place high priority on employers being positive community actors and protecting customer data.

Avoid regulatory scrutiny or sanctions, such as those faced by Apple when apparent bias in the algorithm that sets Apple Card spending limits triggered investigations into the company's use of AI.

Reduce the likelihood of data breaches or misuse. Cases such as the massive customer data theft at retailer Target and, more recently, Meta's apparent violations of EU data regulations have proven these incidents come with punishing costs that can drag on for years.

Generate positive environmental outcomes. Efforts to measure and reduce the carbon intensity of computing and cloud usage through tools like Thoughtworks' open-source Cloud Carbon Footprint (CCF) open the door to aligning technology with the organization's overall sustainability strategy and the path to net zero.

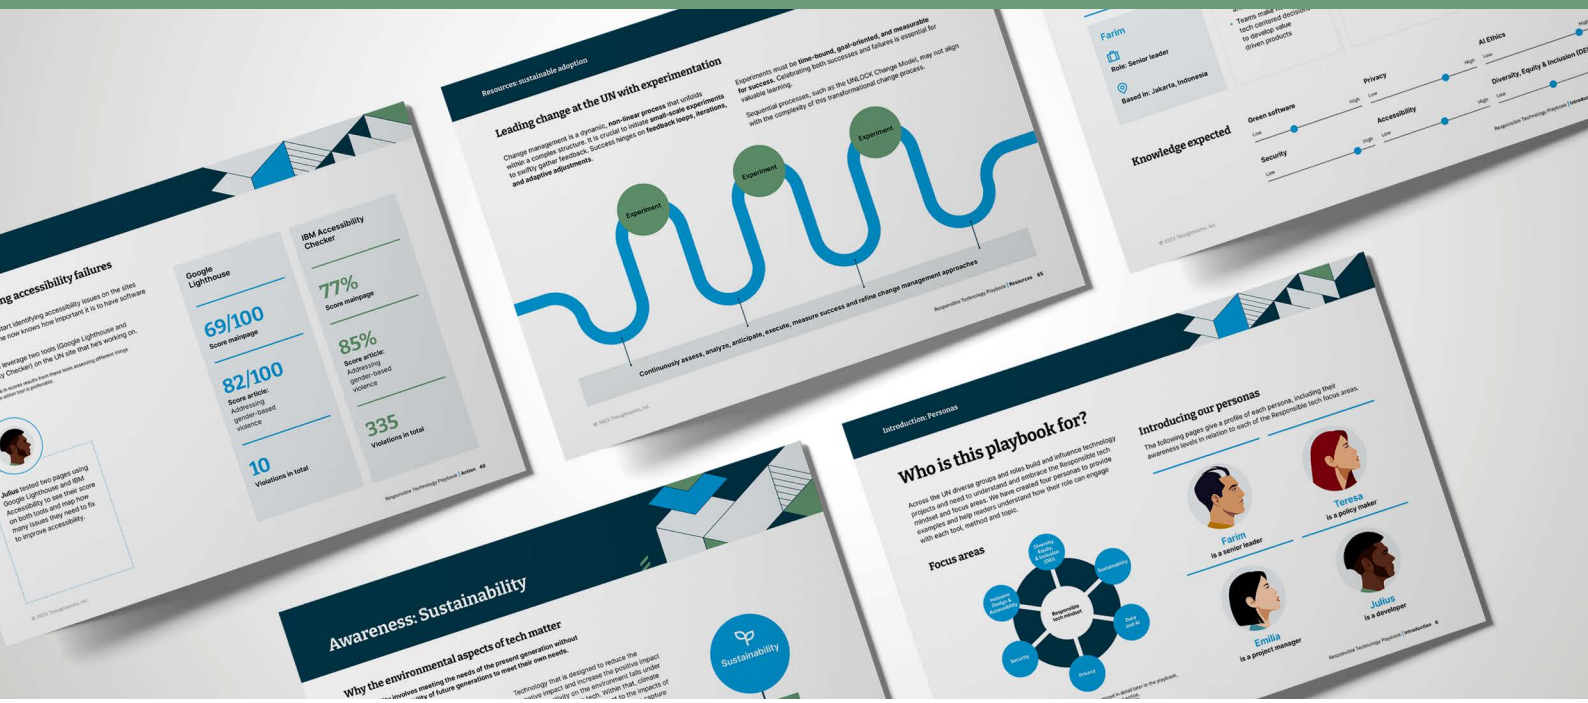

What we've done

Responsible tech playbook with the United Nations

In line with the UN Secretary-General's Strategy On New Technologies, the United Nations Secretariat worked with Thoughtworks to provide guidance on ensuring inclusivity, awareness of bias, transparency and the mitigation of negative unintended consequences in emerging technologies. Following a series of interviews and workshops with United Nations staff, the Thoughtworks and UN team developed a framework and set of approaches for the responsible creation and management of technology systems and products.

"Responsible tech is not just about being aware of what could happen as a result of our well-intentioned actions. It's about being fully engaged with the now — by constantly reevaluating who and what we're protecting and how we're doing it; we're never done."

Rebecca Parsons
Chief Technology Officer, Emerita, Thoughtworks

Actionable advice

Things to do (Adopt)

- **Treat responsible tech practices as a cross-functional requirement.** As mentioned in the title of this lens, including ethical considerations is critical to all organizations.
- **Continuously update technology planning and processes** to incorporate techniques and exercises that help map out the broader consequences of solutions you apply or develop — for example by involving underrepresented groups in design and testing, or simulating breaches that show how data could be misused. Make these techniques part of every process.
- **Establish clear guardrails and policies governing the use of AI** and ensure these are communicated not only to technologists, but to other parts of the organization where more people will be experimenting with AI tools in their day-to-day roles.
- **Adopt secure software delivery practices**, such as making secure development a collective responsibility, producing clean, transparent and easily maintainable code, and continuous testing.
- **Examine your software development processes and tools to understand where you can make more sustainable decisions.** Understand the cost / benefit trade offs of green software engineering techniques.
- **Be a good consumer** by taking steps to understand your partners' and suppliers' stances on responsible technology, and making efforts to engage and support organizations that demonstrate commitment to ethical technology usage.

Things to consider (Analyze)

- **Developing trustworthy data sources** by examining the provenance of information; gathering data where possible from providers that have been vetted and forging partnerships with trusted organizations in your space that govern data-sharing and exchange.
- **Constantly considering changes to what constitutes responsible technology.** Technology is quickly evolving; the problematic activities of tomorrow may not even be possible with the technology of today.
- **Adopting a code of ethics for software development**, either by developing principles that are customized to your organization, or by building on or promoting pioneering standards like the ACM/IEEE-CS Software Engineering Code.
- **Utilizing green software development techniques** such as implementing real-time power consumption monitoring to keep emissions to the minimum viable level, optimizing infrastructure and algorithms and carefully selecting both the location and timing of computation.

Things to watch for (Anticipate)

- **Growing opportunities and threats from developments in AI.** As the frontiers of what is possible for AI to create — or manipulate — rapidly expand, it promises to make significant contributions to everything from market research to product development. However, it will also vastly accelerate the scale and reach of destructive forces like deepfakes and misinformation campaigns. Ensure your organization remains cognizant of and prepared for the new dilemmas AI will present even as you take advantage of its capabilities.
- **Evolving regulations.** You should expect to see regulatory changes across the entire gamut of responsible technology areas. Educate your compliance organizations about new regulatory bodies or agencies which need to be monitored and potentially engaged with across the broad areas we've covered here.

Glossary

A

Addictive tech: Some applications are specifically designed to be addictive, driven by fierce competition for eyeballs and engagement. While this might be good for companies in the business of selling advertisements to audiences, there is an increasing awareness of the societal and environmental harms of addictive tech.

Adversarial machine learning: These are attacks on (or using) machine learning systems. Attackers may tamper with training data or identify specific inputs that a model classifies poorly to deliberately create undesired outcomes.

Affective (emotional) computing: A collective term for systems and devices that can recognize, interpret, process, simulate and respond to human emotions.

Agent-based simulation: The use of simulated independent agents, each working towards their own goals, to model a real world situation. Such simulatioms can help us understand complex phenomena such as the spread of diseases or protein folding.

AGI research: An Artificial General Intelligence (AGI) has broad capabilities across a range of intellectual tasks, and is often compared to human-level intelligence. This contrasts with today's "narrow" AI which can be remarkable, but only for very specific tasks.

AI agents: Functionality built into applications which combines the functionality of publicly available generative AI models with specific knowledge from outside the model, such as product information.

AI as a service: "Ready-to-go" AI solutions offered as a service on cloud platforms. They often don't require specialized AI or ML skills to be used.

AI in security: AI is increasingly being deployed both defensively, to respond to threats more dynamically, and offensively, to probe for weaknesses in a system.

AI marketplaces: Marketplaces such as AWS Marketplace, Google TensorFlow Hub and MS Azure Marketplace enable independent developers and companies to sell their models to a global market. They also allow consumers to quickly leverage those models to create value quickly.

AI safety and regulation: Government regulation and guidance on the use of AI, intended to ensure responsible use and consequences of AI systems. This includes monitoring, compliance and good practice.

AI-assisted software development: The use of AI to speed up or improve software development. Examples include code completion in IDEs, AI-created automated tests, AI that can detect bugs or even AI code generation tools.

AI-generated media: Images, audio or video that have been manipulated by AI. Also known as synthetic media.

AI, IoT and XR combined solutions: A new breed of solutions in which multiple technologies are combined and act together. Drones, robotics and autonomous vehicles are all examples of devices that require machine learning, processing streams of data and layers of intelligence to solve problems.

AI/ML on edge: The ability to run AI and machine learning algorithms at the edge of a network, often on resource-constrained devices.

Alternative currencies: Currencies other than money, such as cryptocurrencies or reputation-based currency. Increasingly, this includes vendor-specific reward-based currencies such as Starbucks Stars or Amazon Coins.

Augmented reality: Where the physical world is combined with the digital. A limited form of AR is now ubiquitous, delivered via Apple and Android mobile devices, capable of overlaying virtual objects to a camera view of the world. More advanced AR is delivered via a dedicated headset such as Apple Vision Pro, Microsoft's Hololens or Meta's Quest 3.

Automated compliance: The use of technology to make all the data required to satisfy compliance reports, checks and balances readily available. In many cases, the automation simplifies reporting by sifting through data; however, AI is now beginning to replace manual decision making.

Automated workforce: The use of technology to perform repeatable or predictable workflows. Automated workforce doesn't mean completely replacing humans; in some cases human-machine "teaming" may produce better results than either working alone.

AutoML: An approach to partly automate the work of data scientists and machine learning engineers by automatically selecting and training machine learning models for specific tasks.

Autonomous robots: Smaller and cheaper than their industrial counterparts, robots with on-board AI are able to sense their environment, navigate, learn to complete tasks and even fix themselves and other things.

Autonomous vehicles: Self-driving cars, trucks and public transport. While the headline focus may be on self-driving cars, autonomous vehicles also have high potential for specialized industrial and business applications such as mining and factory floors.

B

Brain computer interfaces: A device that reads and analyzes signals from the brain and turns them into an input mechanism for a computer. The human and the device, after a period of training, work together to encode and decode human intentions.

C

Causal inference for ML: Techniques to draw cause and effect relationships between the input data and the outcomes of a machine learning model, which allows a model to be more generalizable and require less training data to perform effectively.

Code of ethics for software: A set of guidelines organizations can use to manage risk and mitigate the potential negative consequences of given technologies (such as AI bias).

Collaboration ecosystems: When individuals or organizations share common goals, they will probably want to work together. To do so, they need a set of tools and resources they can use to unlock value effectively — a good example is a remote environment for development teams. It allows people to solve problems together.

Consumer XR: Extended reality intended for consumers rather than professional or enterprise users.

D

Data clean room: Secure environments for organizations to share and combine data with each other without having to physically share their own data

Data contract: A formal agreement between two parties — producer and consumer — to use a dataset or data product.

Data marketplaces: A system that enables the finding, buying, sharing and selling of data within and outside an organization.

Data mesh: A data platform organized around business domains where data is treated as a product, with each data product owned by a team. To enable speed and drive standardization, infrastructure teams provide tools that allow data product teams to self-serve.

Data product specification: A precise technical description of a data product that enables its provisioning, configuration, and governance.

Decentralized data platforms: Use of multiple data stores instead of singular, monolithic centralized stores. A good example is data mesh (see above).

Decentralized identity: Also known as self-sovereign identity, decentralized identity (DiD) is an open-standards-based identity architecture that uses self-owned and independent digital IDs and verifiable credentials to transmit trusted data. Although not dependent on blockchains, many current examples are deployed on them as well as other forms of distributed ledger technology, and private/public key cryptography, it seeks to protect the privacy of and secure online interactions.

Decentralized personal data stores: A data architecture style where individuals control their own data in a decentralized manner, allowing access on a per-usage bases (for example, Solid PODs).

Decentralized security: Rather than using traditional security perimeters that are a single point of failure, techniques such as zero-trust networks decentralize security checks across the network.

Decision science: Combines AI tools and techniques with behavioral and management sciences for the purpose of upskilling and amplifying decision making and decision makers across a variety of complex problems from scenario planning to operations research.

Developer experience platforms: Platforms which provide the tooling to make it as effective as possible for developers to create, test and deploy software.

DevSecOps: An abbreviated portmanteau for development, security and operations. This is an approach that includes security as a first-class concern, together with development and operations.

Differential privacy: A privacy technique that introduces noise in a dataset in such a way as to provide individual privacy while still allowing insights to be drawn or machine learning models to be built on top of the data.

Digital carbon management: Measuring organizational green house gas (GHG) emissions and efforts to mitigate those emissions. Establishing a carbon footprint and a program to determine it is an essential component on the journey towards net zero and is the first building block towards any sustainability strategy.

Digital ecosystems: Disparate participants, systems and even organizations that cooperate, collaborate and compete to create an emergent ecosystem where the whole is greater than the sum of the parts. Examples include the travel industry, online marketplaces and new “super apps” such as Gojek and WeChat.

Digital humans: AI-powered virtual assistants and non-playable characters that recreate human interaction within the metaverse.

Digital twin: A virtual model of a process, product or service that allows both simulation and data analysis. 3D visualization can be used together with live data, so you can understand what is happening to pieces of equipment you can't actually see.

Distributed energy resources: A category of electrical power generation that are "behind-the-meter." DERs generate power for the grid, and reward energy credits to the DER owner. An example is solar panels installed on a home.

E

Easing access to Generative AI: Making AI easier to use by lowering the barrier to entry with shared context and other data that those who aren't familiar with prompt engineering may struggle with.

Edge computing: Bringing data storage and processing closer to the devices where it is stored, rather than relying on a central location that may be thousands of miles away. Benefits include reduced latency for real-time systems and improved data privacy.

Encrypted computation: The ability to perform calculations on encrypted data, without first decrypting it. Useful to maintain data privacy while allowing data storage and manipulation to be outsourced. This includes technologies like secure multi-party computation and homomorphic encryption.

Enterprise XR: An umbrella term for virtual and augmented reality and related technologies which are now being used in the enterprise. Advantages can include cost reductions, efficiency or safety improvements.

Ethical frameworks: Decision-making frameworks that attempt to bring transparency and clarity into the way decisions are made, especially around the use of AI and potential bias in data.

Evolutionary architectures: In contrast to traditional up-front, heavyweight enterprise architectural designs, evolutionary architecture accepts that we cannot predict the future and instead provides a mechanism for guided, incremental change to systems architecture.

Explainable AI: A set of tools and approaches to understand the rationale used by an ML model to reach a conclusion. These tools generally apply to models that are otherwise opaque in their reasoning.

F

Federated learning: An approach that downloads a machine learning model and then computes or trains a specific, modified model using local data on another device. The approach helps multiple organizations to collaborate on model creation without explicitly exchanging protected data.

Fine grained data access controls: More granular access controls for data, such as policy-based (PBAC) or attribute-based (ABAC) that can apply more contextual elements when deciding who has access to data.

FinOps: The practice of bringing financial accountability to the variable spending model of cloud computing. It involves a collaborative approach among teams such as finance, operations and development to manage and optimize cloud costs effectively.

G

GenAI tools in IDEs: The integration of generative artificial intelligence (GenAI) capabilities into integrated development environments (IDEs), the software applications that programmers use to write code.

Generative AI: AI that creates text, image, audio and video from simple human language prompts.

Gesture recognition: Machine understanding and interpretation of human gestures such as waving, making an “up” or “down” motion, hand positioning and so on.

Green cloud: Data centers fed by renewable energy, running software and systems designed and optimized for efficient processing while also minimizing energy consumption.

Green software engineering: Choosing technologies, programming languages, algorithms and software architectures that are efficient and use less energy.

Green UX: Design of user interfaces and prompts that help people understand the environmental consequences of the choices they make. Examples include an airline website displaying carbon emissions for flights or a mapping tool showing the carbon output for driving a particular route.

I

Increased regulation: The steady increase of regulation, especially around data, privacy, security and greenhouse gas emissions.

Industrial XR: Using virtual environments to test and model desired physical outcomes in an industrial context.

Integrated data and AI platforms: Platforms designed specifically for machine learning, providing end-to-end capabilities such as data management, feature engineering, model training, model evaluation, model governance, explainability, AutoML, model versioning, promotion between environments, model serving, model deployment and model monitoring.

Intelligent machine to machine collaboration: Technologies enabling the direct interaction of devices and information sharing between them, usually in an autonomous fashion. This enables to decision making and action with little or no human intervention.

International law for crypto assets: Crypto assets are traded across the world. Similar to the move for international laws for AI, crypto assets also need international law for cross border trading. This might include costs, categories of assets and what constitutes legal trading

K

Knowledge graphs: A way to represent knowledge and semantic relationships between entities using a graph data structure.

M

MLOps: A movement to bring DevOps practices to the field of machine learning. MLOps fosters a culture where people, regardless of title or background, work together to imagine, develop, deploy, operate, monitor and improve machine learning systems in a continuous way. Continuous Delivery for Machine Learning (CD4ML) is Thoughtworks’ approach to implement MLOps end-to-end.

Multimodal AI: AI model interactions that span different modes of communication. For example, a chatbot that understands and responds in both written and spoken language.

N

Natural language processing: Artificial intelligence and other modern technologies that help computers understand the intent and meaning of spoken or written language. Used for everything from dictation software to analyzing documents for meaning.

Next-generation cryptography: Forms of cryptography created in response to technological or societal challenges. Examples include quantum-resistant encryption algorithms, confidential computing with specialized hardware secure enclaves,

homomorphic encryption allowing computation to occur on the data while it is still encrypted, and energy efficient cryptography.

O

Online machine learning: A technique where algorithms continuously learn based on the sequential arrival of data, and can explore a problem space in real time. Contrasts with traditional machine learning where model training uses only historical data and cannot respond to dynamic or previously-unseen situations.

Operationalize AI: Making AI a normal part of business operations including appropriate security and governance.

P

Personal information economy: A business model that aims to extract business value from the possession and use of large amounts of personal information. Examples range from the primitive use of cookies to the targeted profiling of people via their online behavior. This has historically been the domain of companies or intermediate ad-based services trying to retain and target customers, but, since GDPR and similar privacy laws, we are seeing a shift towards people controlling what data they wish to expose in exchange for a service.

Personalized healthcare: Understanding an individual patient's genetic profile to identify potential issues before they happen and provide more effective treatments in response to existing conditions.

Platforms as products: A way of creating and supporting platforms with a focus on providing customer (user) value instead of treating platform building as a time-boxed project.

Privacy first: Privacy first is a significant shift in business, organization and product strategy, where privacy operates as a core business value and offering. This shift moves away from the prior movement where "users are the product", into a new realm, where building trust and transparency comes first.

Privacy-aware communication: Communications software that directly advertises its security stance and features, such as end-to-end encryption.

Privacy-enhancing technologies (PETs): A collection of technologies and techniques for preserving user privacy, such as anonymization, encrypted computing and differential privacy.

Privacy-respecting computation: New techniques that allow stronger guarantees for privacy, even when personal data is used in computations. Part of the broader category of privacy-enhancing technologies (PETs).

Production immune systems: Systems that monitor metrics across complex distributed systems and take corrective action if a problem is detected. They are often used for security, but increasingly also for resilience and recovery in the face of an outage.

Q

Quantum computing: Use of probabilistic states of photons, rather than binary ones and zeros, to run algorithms. Although proven to work in specific problem spaces, quantum computing has yet to scale to broadly useful applications.

Quantum machine learning: Machine learning algorithms adapted and executed on a quantum computing engine, generally used to analyze classical (non-quantum) data.

R

RAG (Retrieval Augmented Generation): A method in artificial intelligence where the system enhances its response generation by fetching relevant information from a large database or knowledge source. This approach combines the creative aspects of

generative AI models with the precision of data retrieval, enabling more accurate and contextually relevant responses in various business applications.

Re-decentralization: Systems, both human and machine, originally designed to be decentralized have become more centralized over time. Re-decentralization refers to the conscious effort of moving those systems back to a decentralized model.

Responsible tech facilitation: Tools and techniques are emerging that support incorporating responsible tech into software delivery processes, primarily focusing on actively seeking to incorporate under-represented perspectives; some examples include Tarot Cards of Tech, Consequence Scanning, and Agile Threat Modeling.

Retina resolution XR: Ultra-high resolution XR with photorealistic rendering over a wide field of view. Currently only available via extremely expensive headsets. An example is Varjo XR-3.

Robotic process automation and low code: Robotic process automation (RPA) aims to allow scripts or bots to interact with UIs instead of needing a human operator. Low-code seeks to democratize programming, by allowing non-programmers to create software systems.

S

Satellite networks: High-speed, low-latency broadband for places where traditional fiber or wireless network providers won't spend the money to connect. Examples include Starlink from SpaceX, Kuiper from Amazon, OneWeb and Telesat.

Secure software delivery: Security applied to the entire process of software creation, which in modern architectures includes the delivery pipeline used to build, test and deploy applications and infrastructure.

Smart cities: An urban area that uses different types of IoT sensors to collect data coupled with platforms to integrate and act on the data, advising or commanding digitally enabled systems to perform some response. Insights gained from the data are used to manage assets, resources and services efficiently; in return, that data is used to improve the operations across the city.

Smart energy management systems: Ubiquitous availability of energy usage data via measurement equipment, APIs and tools gives a range of energy players (generators, distributors, suppliers, vendors) and customers a greater ability to understand and analyze their energy usage.

Smart homes: Featuring smart hubs, homes are now becoming 'smart', allowing people to control almost all household systems. Analytics can even guide or manage heat and energy supply and learn from individual habits or those in a neighbourhood.

Smart systems and ecosystems: Networks of networks that use AI and ML to enhance a system to become more than the sum of its parts. For example, in a smart city, networks of cars and roadside sensors help speed the flow and safety of traffic.

Software-defined vehicles: Automobiles where the core functionalities, features and user experience are primarily governed by software, rather than traditional mechanical and electrical systems. This approach enables increased flexibility, customization and continuous enhancement through remote updates, significantly transforming the vehicle's capabilities and, in turn, the automotive industry's business models.

Spatial audio: Advanced signal processing, originally from Apple, that allows sounds to be placed virtually in 3D space. Spatial audio also tracks headphones and screen position to allow for accurate sound placement.

T

Technology and sovereign power: Rising forces are leading to internet balkanization — the splintering of the internet — many led by nation states. Privacy legislation accelerates this process, as it enforces data rights, data sovereignty, and strongly impacts how companies deploy and distribute systems and data on the Internet.

Technology for circular economy: A closed economic system where raw materials and products are constantly shared so as to lose their value as little as possible. Technology that supports this includes reusable services, traceability, IoT and data mining.

Touchless interactions: The ability to interact with devices without touching, driven at least partially as a result of the COVID-19 pandemic. Specific technologies include hand tracking and voice and gesture recognition.

Trustworthy data: An emerging set of techniques to certify the provenance of data and to govern its use across an organization. This could prove transformative in the effort to track and enhance progress towards sustainability targets.

U

Ubiquitous connectivity: Providing connectivity to everyone and everything, everywhere, all the time. Some predict ubiquitous connectivity will super-charge innovation in resource-limited parts of the planet, while critics see it as expensive and unnecessary.

Understandable consent: Most terms of service (TOS) or end-user license agreements (EULAs) are impenetrable legalese that make it difficult for people without a law background to understand. Understandable consent seeks to reverse this pattern, with easy-to-understand terms and clear descriptions of how customers' data will be used.

V

Vector databases: Specialized storage systems designed to efficiently handle and index high-dimensional data vectors, commonly used in machine learning and AI applications.

X

XR-enabled hybrid working: A collaboration strategy where, using XR, everyone on a hybrid local/remote team interacts with the same shared artifacts, such as whiteboards and other information radiators. This brings the remote collaborators closer to the in-person team.

Z

Zero knowledge proofs: A method that allows one party to prove to another that a statement is true without revealing how it knows it is true.

Thoughtworks is a global technology consultancy that integrates strategy, design and engineering to drive digital innovation. We are over 11,000 Thoughtworkers strong across 51 offices in 18 countries. For 30 years, we've delivered extraordinary impact together with our clients by helping them solve complex business problems with technology as the differentiator.

/thoughtworks

Strategy. Design. Engineering.