

Looking Glass

Bringing tech-led business changes into focus

/thoughtworks

Strategy. Design. Engineering.

2025

Introduction	3
Operationalizing AI for business impact	4
Strengthening the data value chain	11
Reimagining responsible tech for the era of generative AI	18
Enabling richer experiences through multimodal interactions	26
Unlocking greater value from physical-digital convergence	34
Glossary	42



Introduction


Welcome to the Looking Glass 2025. Unlike many tech trend reports, Thoughtworks' Looking Glass is not intended to shine a light on the latest buzzwords. Instead, we take a long term look at the technology horizons and explore what that means for businesses. What are the things you need to know about now? And what's likely to be important in the longer term? The Looking Glass enables you to understand and interpret emerging technologies so you can make sound, strategic choices for your organization.


The relentless speed of technological advancement makes it harder to predict what's coming and where your investments will pay off the most. Breakthroughs in areas such as agentic AI promise to upend how we think about technology. But how quickly should you prepare to adapt? Here's where Thoughtworks' Looking Glass comes in.

In this edition, we explore more than 90 trends through five distinct perspectives that define the evolving tech landscape in business. Some of these trends are already transforming operations, while others remain just over the horizon, sparking interest and debate but still unfolding. For business leaders, keeping a broad, strategic perspective on these developments — both current and future — is essential. Looking Glass offers exactly that: a framework to gain a comprehensive understanding of key trends.

The five lenses provide clarity and focus, helping ensure your organization remains adaptable, resilient and ready to harness or respond to the inevitable shifts in technology that shape our modern world.

—
Rachel Laycock
Chief Technology Officer, Thoughtworks





Operationalizing AI for business impact

The mainstreaming of AI — and generative AI in particular — is continuing apace. But as AI proliferates, it's more evident that successfully operationalizing AI models and bringing them to production remains a challenge. From questionable output to unintended consequences, there are a host of real and projected scenarios that prevent organizations from leveraging AI to its full potential.

Enterprises continue to struggle with data quality, data accessibility and the challenges of data at scale, all of which remain foundational to robust, effective AI. As our data platform lens explores, careful data curation, and effective data engineering and architecture are essential. The importance of synthetic data, particularly in research contexts, as a tool to avoid privacy and data integrity issues is also becoming more and more apparent.

Organizations also need to develop better approaches to the evaluation and control of AI systems. Forward-looking enterprises are adopting 'evals' — tests of AI output to determine reliability, accuracy and relevance — and guardrails, programmed policy layers that mitigate the inherent unpredictability of generative systems.

As adoption increases, improving the mechanisms through which AI systems are connected with enterprise applications grows more important. Proxy services are emerging to help developers link AI models with the applications they build.



AI agents are sometimes positioned as the next step in the evolution of AI, due to their capacity to mimic human reasoning. However, the technology remains relatively new, and finding applications for agents requires domain expertise, as well as the ability to precisely map and model complex processes and interactions. To build a sustainable and productive AI practice, it's vital that the organization doesn't resort to shortcuts, acquires the requisite skills and keeps innovation rooted in business realities.

“The lessons from automation endeavors in the ‘80s could help to build the right level of human-AI agent handovers. We must focus on augmenting humans rather than trying to substitute their current tasks completely.”

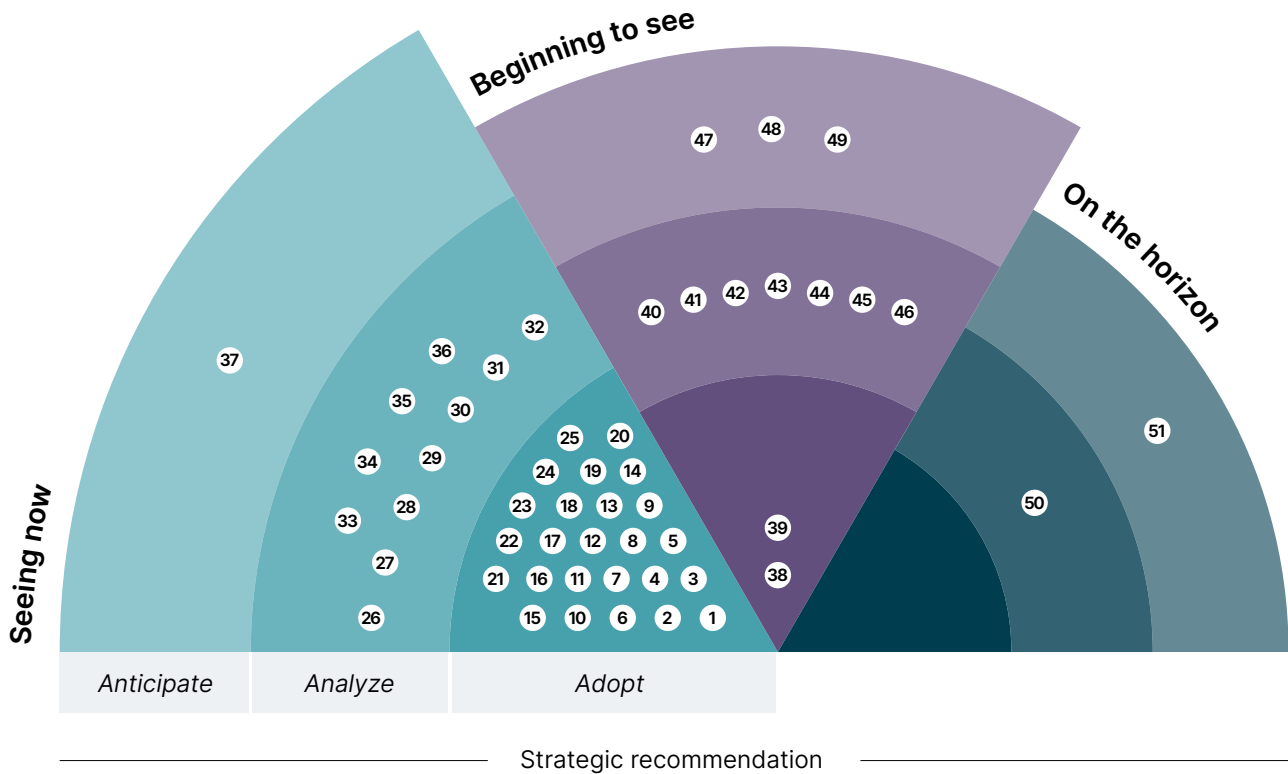
Srinivasan Raguraman
Technical Principal, Thoughtworks



Signals

- **The emergence of small language models**, such as [Microsoft's phi-3](#), and [AMD's AMD 135](#). These make it possible to run AI models at the edge of networks on devices like mobile phones, and because they are relatively lightweight, focused and efficient, have a range of positive [business, security and sustainability implications](#). LLMs also continue to evolve, with Anthropic's Claude 3.5 Sonnet LLM, which has set industry benchmarks in terms of performance, recently upgraded to include [computer use capabilities](#).
- **Research showing that for many organizations, AI investments and adoption aren't necessarily translating into deployment or business impact**. While interest in (and spending on) AI solutions remains high, businesses are beginning to [pay more attention to the cost of AI projects](#), and stepping up efforts to ensure they [deliver value](#).
- **The coming into force of the [European Union's AI Act](#)**, which sets an international benchmark by laying out [obligations around data governance, documentation, human oversight and security](#) for businesses adopting AI systems.
- **Sustained, massive investment in data centers**, with [Google even turning to nuclear power](#) to generate the vast amounts of power its AI offerings are likely to require. This indicates AI is a long-term bet that will continue to gain momentum in the business context, and in society as a whole.
- **The growth of tools simplifying how engineers and others interface with AI models**, such as [LiteLLM](#) and [Langchain](#).
- **Renewed focus on tackling LLM 'hallucinations' and fabrications**, with novel techniques like ['semantic entropy'](#) being applied to root out errors, and LLMs policing the output of other LLMs.
- **Rising awareness of 'shadow AI'**, or the use of unsanctioned AI tools in the enterprise context, which could pose significant problems for companies if sensitive information is leaked to LLMs by employees. In [one recent survey](#) a third of organizations admitted to finding it hard to monitor the illicit use of AI among their teams.

Trends to watch



Seeing now

Adopt

1. Accessibility in multimodal experiences
2. Agent-based simulation
3. AI agents
4. AI as a service
5. AI in security
6. AI-assisted software development
7. Automated compliance
8. Collaboration ecosystems
9. Data mesh
10. Edge computing
11. Ethical frameworks
12. Evaluating and managing AI outputs
13. Evolutionary architectures
14. Explainable AI
15. Generative AI
16. Integrated data and AI platforms
17. Interfacing with AI
18. LLMOps
19. MLOps
20. Model training optimization
21. Online machine learning

22. Platforms as products
23. Privacy first
24. Software-defined vehicles
25. Vector databases

Analyze

26. AI marketplaces
27. AI safety and regulation
28. AI-generated media
29. Automated workforce
30. Autonomous robots
31. Changing perceptions of AI
32. Easing access to generative AI
33. Federated learning
34. Multimodal AI
35. Personalized healthcare
36. Synthetic data

Anticipate

37. Understandable consent

Beginning to see

Adopt

38. AI-ready data
39. Fine grained data access controls

Analyze

40. AI Observability
41. Data lineage
42. GenAI computer control
43. Intelligent machine to machine collaboration
44. Production immune systems
45. Small language models
46. Talk to data

Anticipate

47. Adversarial machine learning
48. Affective (emotional) computing
49. AI in robotics

On the horizon

Adopt

—

Analyze

50. AI avatars

Anticipate

51. AGI research

The opportunities

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



Enhance knowledge management and transfer by adopting GenAI to help employees sift through, summarize and analyze stores of enterprise data, whether structured or unstructured. A wide range of products are emerging to facilitate the retrieval and dissemination of important information in industries like property.



Harness AI to accelerate processes like legacy modernization and coding. Thoughtworks is already successfully applying GenAI to assist teams with one of the most difficult aspects of modernization: understanding and unpacking the intricate web of connections that typically underpin legacy systems and codebases. AI assistants can also significantly boost the productivity of software development and other teams by taking over frequent, repetitive tasks.



Explore AI agents to elevate automation, potentially transforming how employees perform tasks like scheduling and customer support, and raising the bar for engagement and personalization in customer interactions.



Boost the speed at which LLMs are brought into production, and their effectiveness when deployed through emerging practices and tools like LLMOps, which accelerate model development; retrieval-augmented generation (RAG), which can enhance models' reliability; and AI gateways or smart endpoints to connect AI systems to applications.



Develop and communicate a joined-up AI strategy that empowers employees to experiment with AI in a structured way, while preventing the emergence of 'shadow AI' that could pose a threat to the organization's intellectual property or reputation.



Leverage small language models to bring AI innovations to edge devices, offering opportunities for everything from operational analytics to personalization — without compromising privacy, since data doesn't have to be moved to the center of a network.



Lead the way in terms of compliance and ethical AI practices. We urge our clients not just to follow but embrace regulations like the EU AI Act, as such legislation often reflects wider societal sentiment and concerns — and potential customers take notice of businesses that are responding.



What we've done

PEXA

Thoughtworks partnered with digital property technology company PEXA, AWS and Redactive to develop an innovative and versatile AI assistant that has boosted the productivity of PEXA's employees by providing personalized answers to queries and augmenting tasks like information discovery.

Seamlessly integrated with PEXA's internal systems, the solution also met robust requirements for data security and privacy by equipping the assistant with permissions awareness, ensuring employees are only able to access information cleared for sharing.

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.
- **Implement a holistic and comprehensive AI strategy** for your organization that includes guidelines on permitted tools and the contexts in which AI can be used, to minimize the risks of shadow AI.
- **Adopt retrieval-augmented generation (RAG) when developing AI systems**, to give reliability an uplift and position models to create more specific outputs. Integrating evals and observability can further enhance the resilience of systems over the long term.
- **Embed AI throughout the software development lifecycle.** Maximum results are achieved when the role of AI isn’t just limited to coding, but assists with processes like testing and documentation.
- **Apply data mesh and data product thinking** to ensure AI applications are built on the robust data foundation needed to ensure they deliver business or customer value. Disciplines like data curation, which creates, organizes and manages data sets so they’re transparent and easily accessible, also contribute to the success of AI.
- **Use proxies to simplify the way teams interact and leverage AI models**, paving the way for the enhancement of applications they develop with AI features and capabilities.




Things to consider (Analyze)

- **Avoid what's known as the 'substitution myth'** — the idea that AI can simply directly replace a human. Instead, build and implement systems that augment roles to make teams more productive and engaged, while acknowledging the continued importance of human judgement and oversight.
- **Be cognizant of varied expectations around AI.** Research suggests people may approach AI differently depending on cultural background, with some wanting a high degree of control and others prioritizing a sense of connection. These differences, as well as variances in context or situation, need to be understood and acknowledged when planning and implementing AI.
- **Pay close attention to costs**, and try to identify the approaches most likely to meet your needs while generating return on investment. Running AI models can be expensive, especially if expenses like employee compensation are factored in. Keeping spending in check requires active financial monitoring (i.e. FinOps) and consideration of things like small language models.
- **Monitor AI regulation and future policy developments**, particularly how these intersect with privacy laws, which could have a massive impact on the data resources available for AI projects. Multiple US states, and countries from Canada to India and Japan, are planning to enhance or roll out legislation that will set guardrails around AI use and development.

Things to watch for (Anticipate)

- **Questions around legal liability and accountability** for the negative consequences of AI use. As issues such as AI misleading customers and the associated legal challenges emerge, authorities like the EU are moving to make organizations more culpable.
- **The potential growth of AI companions**, designed to provide emotional support, friendship or even intimacy. While these could help combat loneliness and isolation, they may also have troubling implications for human interaction, requiring businesses to think carefully about the introduction of AI with companion-like features.



Strengthening the data value chain

Leveraging data platforms and AI

As enterprise adoption of AI gains pace, there's rising awareness of data's role as a differentiator, and a source of competitive edge. Developing the capabilities to leverage data at speed and scale, and become truly data-driven, has become an emerging priority. Treating data as a product represents one of the most effective means to achieve this goal, and the best way to build and distribute data products is through data platforms.

The principles that underpin high-performance data platforms remain the same — decentralization and federated data ownership — but new trends and opportunities in the space are presenting challenges that organizations need to be prepared for. In particular, the rise of generative AI (GenAI), and the importance of unstructured data in it, requires teams to think differently about how data is managed and processed. It's becoming critical to treat unstructured data as a first class citizen, not as structured data's poorer cousin.

It's also important to note the rising need for better — and ideally automated — governance of data products.

Data products — reusable data assets engineered to deliver trusted datasets for specific purposes — exist in dynamic environments where the needs of teams and the wider organization are constantly evolving, and it's important that they also develop in a way that delivers value.



Maintaining the capacity for competitive and sustainable change requires intentional design of cohesive centralized and decentralized capabilities. Some organizations are navigating away from creating consensus-based 'single sources of truth' to forming integrated 'contextual truths'.

Equally essential is ensuring data products are built with a clear line to business adoption. Platform and product thinking can help, but there's a need to move beyond existing paradigms and tooling, and consider applying human-centered design for more effective ways for data to be consumed and leveraged by business users. GenAI and trends like 'talk to data' and graph-based discovery are creating promising opportunities in this space, transforming the way teams interact with and consume data.

“An open and evolving data and AI platform allows organizations to embrace uncertainty in rhythm with changing demands, fostering a culture of continuous learning.”


Nimisha Asthagiri
Technical Principal and Data Mesh Leader, Thoughtworks



Signals

- **Unstructured data moving from a supporting to a starring role.** There's growing focus on the use of unstructured data (such as text, video, images and audio) to build better AI training models, which requires integrating and working across different types of data in as frictionless a way as possible. Startups in this space are gaining significant investment and the likes of IBM are unveiling new products specifically designed to help enterprises unleash the potential of unstructured data in analytics and AI.
- **Enterprises applying GenAI to better leverage unstructured data.** GenAI's ability to parse and summarize vast quantities of the information contained in everything from meeting recordings to PowerPoint presentations, and to support natural language interactions, is transforming the way teams access and use data and enhancing knowledge management. However, this trend is also raising questions as to whether AI and GenAI platforms should be integrated with other data platforms or kept distinct, which, in some cases, is leading to platform proliferation.
- **More organizations grappling with the challenges of treating data as a product,** as it becomes a business imperative. Research shows the vast majority of businesses see clear benefits from such an approach, including improved data sharing and strengthening the connection between data and business goals. However, they are confronting multiple barriers along the way, from fragmented systems to uncertainty about data provenance.
- **The rising importance of data discoverability.** By empowering users to better discover, understand and use data assets, data catalogs can play an important role in data platforms and a data product approach. But they can also cause more issues than they solve if their user experiences or capabilities are limited, impeding the discovery process. The recent introduction of knowledge graphs to data platforms is addressing these risks, making it possible to draw out relationships and nuances in data that are typically lost in the process of abstraction.
- **More pressure being put on data teams to demonstrate ROI and manage costs more effectively.** The increasingly established link between data strategy and enterprise performance also means these teams can no longer work in isolation; instead strategies should be co-developed with, and create platforms that deliver results for, the business.

Trends to watch



Seeing now

Adopt

1. AI as a service
2. Automated compliance
3. Collaboration ecosystems
4. Data catalog
5. Data fitness functions
6. Data mesh
7. Data product specification
8. Developer experience platforms
9. Digital twin
10. Edge computing
11. Ethical frameworks
12. Explainable AI
13. FinOps
14. Green computing
15. Integrated data and AI platforms
16. Knowledge graphs
17. MLOps
18. Model training optimization
19. Online machine learning
20. Platforms as products
21. Privacy first
22. Privacy-enhancing technologies (PETs)

23. Secure software delivery
24. Smart systems and ecosystems
25. Vector databases

Analyze

26. Autonomous robots
27. Autonomous vehicles
28. Data clean room
29. Data marketplaces
30. Decentralized data architectures
31. Federated learning
32. Semantic representational technologies
33. Synthetic data

Anticipate

34. Understandable consent

Beginning to see

Adopt

35. AI-ready data
36. Data contract
37. Fine grained data access controls

Analyze

38. Data lineage
39. Integrating unstructured data
40. Intelligent machine to machine collaboration
41. Talk to data

Anticipate

42. Decentralized personal data stores

On the horizon

Adopt

—

Analyze

—

Anticipate

—

The opportunities

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



Consolidate data and AI platform capabilities, enabling AI as a service to embed this new technology and empower users to leverage it successfully throughout the organization. Surveys have shown that despite concerns about the wider impacts of AI, adoption has positive implications for teams' collaboration, efficiency and performance.



Use AI (and GenAI) to build and maintain data products more effectively. Emerging AI tools have the potential to contribute to data products in a number of ways, from synthesizing and analyzing information garnered in end-user research or testing, to accelerating coding and creating documentation that can smooth the path to effective adoption.



Enhance control over costs. With data management often dominating enterprise technology spending, introducing new tooling to track data lineage and analyze the impact of complex data initiatives can help teams determine and demonstrate ROI with greater precision. FinOps thinking can contribute significantly to this process by strengthening the links between tech and business teams and ensuring investments come with financial accountability.



Strengthen data governance by introducing emerging best practices and structures. These include data clean rooms, secure, self-contained environments where enterprises can blend proprietary and third-party data to improve analytics and personalization while protecting customer privacy; and data contracts, which by setting ground rules for data users and consumers, can improve transparency and trust when sharing data across an organization.



Combine knowledge graphs and GenAI, which can enhance understanding of large, complex data sets by mapping the relationships among entities within them. This opens the possibility of more semantic approaches to integration, which in turn can help create a better user experience for data consumers. In addition, combining knowledge graphs and GenAI can also deliver better LLM responses because we're taking explicit knowledge from knowledge graphs and combining it with implicit statistical knowledge from LLMs.



What we've done

Pfizer

Thoughtworks is working actively with these leading pharmaceutical companies to create data mesh platforms that enhance their ability to create and deliver transformative data products. With Pfizer, we helped develop cutting-edge layered platforms serving AI-powered data products, graph-based semantic interoperability, and LLM-based agents that drive the firm's oncology research, supporting early drug discovery.

Gilead

For Gilead, we supported the design and implementation of Gilead DnA, a scalable enterprise-wide data platform that provides data engineers and researchers with a secure self-service environment for data processing, complete with 'talk to data' functionality.

Actionable advice



Things to do (Adopt)

- **Lay the right foundations** for creating effective data products by implementing a data mesh, which places data within the reach of teams that need it most and reduces friction between data producers and consumers.
- **Automate data governance** as much as possible to ensure policies are implemented consistently and with minimal impact on data usage and consumer experience. Fitness functions and more rigorous monitoring of service level indicators (SLIs) can be good places to start.
- **Start treating unstructured data as a first class citizen** that is given the same attention and prominence as structured data in your data platform, and draw on its potential to improve analytics and AI models.
- **Invest in a superior data product development experience** to accelerate adoption. Mapping decision journeys can help the organization better understand and trace how to move from use cases to data, and particularly AI data, products.



Things to consider (Analyze)

- **Extend user experience and human-centered design to data and AI.** This includes thinking carefully about how to build the best possible interface and experience for discovering and accessing data, out of an expanding range of GenAI-enabled options.
- **Examine ways to track and document data lineage and improve metadata for data products** for data consumers. Doing so can also enhance data governance and data engineering by highlighting opportunities to smooth the flow of data throughout the organization. AI tools can play a valuable role in this process by providing a quick and precise snapshot of data's history and transformations.
- **Adopt mechanisms to minimize the risk of creeping centralization.** Encourage teams to think less about creating a single source of truth and more about adopting federated data management that efficiently delivers what the use case or context demands.
- **Track ROI for data and AI transformations.** It's important to be able to demonstrate the value and impact being driven by data and AI initiatives. There's no single way of doing this, but it's a valuable step in ensuring teams remain value-focused and that projects in this area have organizational buy-in.

Things to watch for (Anticipate)

- **Next-gen user experiences like voice and VR impacting data discovery.** By allowing users to query data naturally and moving data visualization into a three-dimensional space, new tools promise to transform the way teams perceive, interact with and understand information, paving the way for deeper analysis and collaboration.
- **Propagating more granular access controls** as data platforms and products scale to more users and data product development accelerates. Studies show data professionals are already walking a fine line between prioritizing security and not impeding the efficiency and flexibility data platforms are designed to provide.
- **Adopting GenAI and knowledge graphs** to improve data discovery and better describe and document entities in large data sets.




Reimagining responsible tech for the era of generative AI

The importance of ‘responsible tech’ has been widely discussed for a number of years now. However, across the industry it has struggled to properly cut through, remaining a somewhat marginal concern. While this might not be unsurprising during a period of economic uncertainty and tighter budgets, the rise of generative AI has made the topic more urgent than ever. This is because the ethical, legal and even philosophical questions raised by the technology can, at least in part, be addressed through responsible technology principles and practices.

This means 2025 is the year businesses need to properly embrace responsible technology. Without it, attempts to experiment and innovate with generative AI and associated technologies may contain risks that businesses simply do not need during a challenging period, ranging from the financial consequences of compliance failures to damaged consumer trust.

But what does embracing responsible tech in the generative AI era actually involve?

We see it as beginning with a recognition — that has long been absent — that responsible tech isn’t something you bolt on to existing activities and projects: it’s something that needs to be embedded in organizational values, team practices and cultures.



It means considering all the potentially negative consequences of, for instance, a new generative AI chatbot, whether in terms of data leaks and privacy breaches, all the way to disturbing and harmful content being served to end users. It also means anticipating regulatory and legislative demands rather than simply reacting to new legislation; given legislation is often driven by the concerns and interests of the wider public — ie. your consumers — it's invariably a useful heuristic for building trust with those that matter most to your business.

You may need to fight for responsible tech on two fronts: in how these technologies are built and deployed; and in how these technologies are shaping what we do. Generative AI capabilities have been added to all manner of products, which may catch some consumers unaware — for instance, they might not realize the service representative they're chatting to isn't human. This means part of responsible tech isn't just about being decisive — it's also about being sensitive to the unknowns that are inherent in an environment where generative AI is everywhere. This mindset should be extended across the breadth of what your organization does, regardless of whether AI is a priority for your business or completely outside your operational scope.

“Responsible technology isn't a checklist — it's a mindset. In the age of generative AI, embedding ethics at the core of innovation isn't just about avoiding risks; it's about building trust, anticipating change and leading with purpose in a fast-changing world.”

Ken Mugrage
Principal Technologist, Thoughtworks



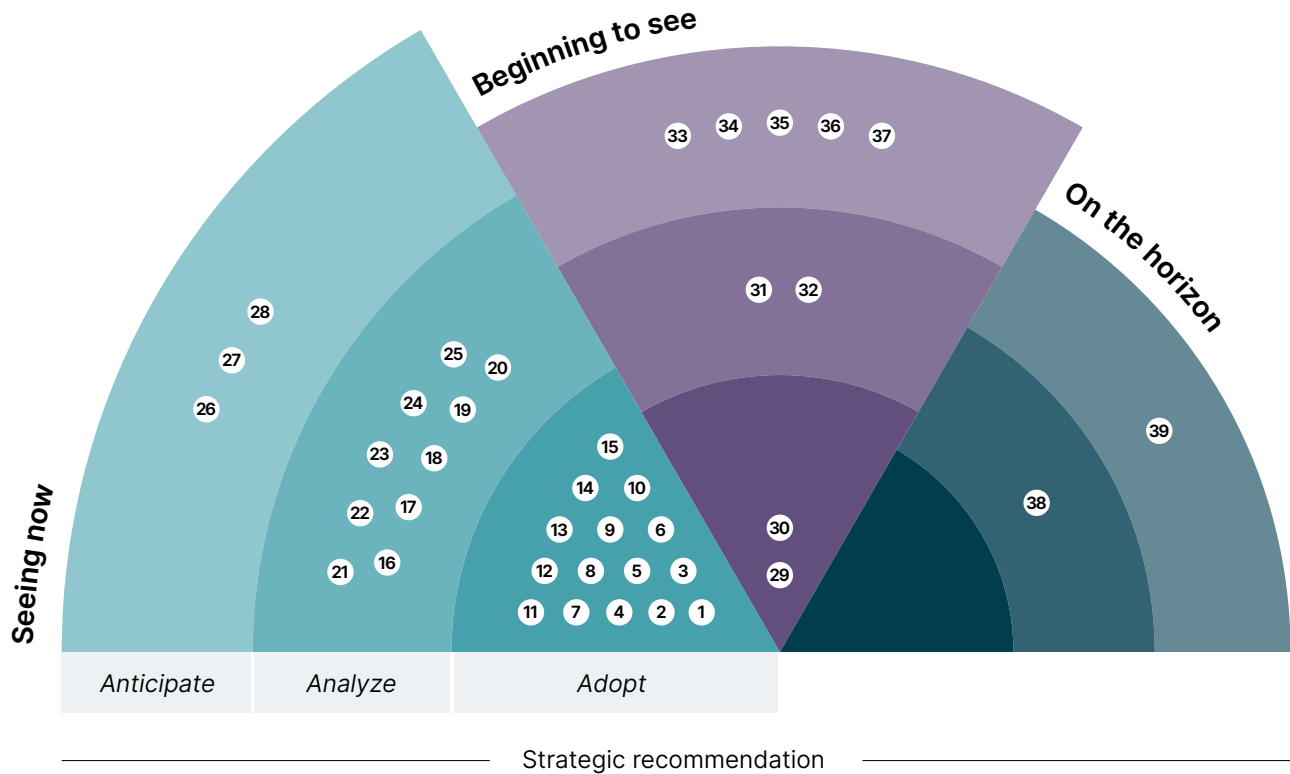
Signals

- **Questions about accountability and legal liability for harmful technology consequences.** The lawsuit of a mother of a teenager who died after interacting with an AI chatbot could be a significant moment in AI legislation and shape how we think about liability and responsibility.
- **Increasing awareness — or perhaps chaos — around what data has and has not been used to train major AI models.** For instance, there was a remarkable moment when an OpenAI leader didn't know whether Sora was trained on YouTube videos, while abusive material was found in the LAION 5-b dataset. There was also confusion when LinkedIn suspended data processing in the UK in September 2024 following concern from the Information Commissioners Office about the way UK users' data would be used to train generative AI.
- **Water consumption of data centers is causing significant concern.** At a local level in areas affected by drought there is even greater political friction, highlighting the ongoing environmental questions raised by AI usage.

- **The consequences of corporate greenwashing** are becoming tangible as organizations are held to account for false claims about their green credentials.
- **The interdependence of the industry** — underlining the importance of trust and transparency at a technical level — demonstrated by the CrowdStrike outage and other supply chain vulnerabilities like the XZ Utils backdoor.
- **Increasing shadow AI in organizations**. As the generative AI market has grown, it is incredibly easy for employees to experiment with AI without oversight. This can create significant privacy risks.
- **The fragmentation of the social media landscape**. The lack of stability in this space — best demonstrated in the mass exodus of X (formerly Twitter) users — underlines significant content and safety problems with platforms as well as growing consumer fears about online safety, privacy, mis/disinformation and even their digital consumption habits more broadly.
- **The growth of impact investing**. This is where investors target big social problems — like health or the environment — with a view to capturing value. It has been called out as something that rather than tackling social issues can actually exacerbate them.
- **Legislation taking on manipulative design**. The EU, for instance, has introduced a law aimed at tackling dark patterns.



Trends to watch



Seeing now

Adopt

1. Accessibility in multimodal experiences
2. AI in security
3. AI-assisted software development
4. Automated compliance
5. Biometric authentication
6. Decentralized security
7. DevSecOps
8. Digital carbon management
9. Edge computing
10. Ethical frameworks
11. Green computing
12. Model training optimization
13. Privacy first
14. Privacy-enhancing technologies (PETs)
15. Secure software delivery

Analyze

16. AI safety and regulation
17. AI-generated media
18. Alternative currencies
19. Automated workforce
20. Changing perceptions of AI
21. Data clean room
22. Data marketplaces
23. Impact funds
24. Internet regulation
25. Synthetic data

Anticipate

26. Addictive tech
27. Next-generation cryptography
28. Understandable consent

Beginning to see

Adopt

29. Data contract
30. Fine grained data access controls

Analyze

31. Data lineage
32. Production immune systems

Anticipate

33. Adversarial machine learning
34. Affective (emotional) computing
35. Decentralized personal data stores
36. Quantum computing
37. Responsible tech facilitation

On the horizon

Adopt

—

Analyze

38. AI avatars

Anticipate

39. AGI research

The opportunities

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



Strengthen organizational culture by working to ensure there is genuine and authentic alignment between corporate rhetoric and values and the perspectives and values of employees. This is certainly not easy at a commercially challenging time, but mistrust and cynicism will have long-term consequences that may prove even more difficult to repair.



Revisit and strengthen those values. Values need to be maintained and evolved — if left without stewardship they will prove useless. Organizations, and leaders in particular, should spend time considering existing values and whether they're relevant and, most importantly, actionable. They need to be things that can guide behaviors and decision-making at every level; they should be things that can be put into practice.



Leverage AI thoughtfully. AI can offer a competitive advantage: but simply rushing to integrate it is risky. The real opportunity is to be thoughtful about why, where and how AI capabilities are used. This will not only minimize potential risks, it will also strengthen your relationship with customers and ensure you are delivering even more value.



Build trust with consumers. Concerns about privacy and the way data is used and managed continue to grow. Businesses that seek to buck the perceived trend for ever-increasing extraction may be able to gain an advantage in the market. Transparency and trust can be a differentiator.



Be more intentional and considered about what you're doing to meet user needs. Foregrounding accessibility practices in your organization can help ensure you are building products, services and experiences that provide value for even more people.



Reduce waste and improve efficiency. While businesses should, of course, be focusing on their environmental impact, in challenging times the bottom line takes absolute priority above everything else. However, it's possible to do both: in fact, framing environmental action in terms of efficiency can be an effective way to ensure it is taken seriously across the business. In other words, business leaders need to present responsibility and commercial impact as things that are closely intertwined, not mutually exclusive.



Focus on skill development. Avoid the temptation to automate everything and rely on AI tools to 'do more with less'. What happens when you don't have knowledge or skills needed to solve problems down the line? Considering how human skills and AI capabilities can complement one another will ensure you have a team that is able to help the organization reach its future objectives.



What we've done

Swann Security

Swann Security partnered with Thoughtworks to develop the world's first AI security concierge, a groundbreaking solution designed to enhance home protection while safeguarding privacy. Leveraging generative AI, the system engages visitors naturally, whether homeowners are present or not. Prioritizing privacy, Thoughtworks crafted a prompt engineering strategy to ensure the AI's responses respect security boundaries and defend against intrusive or adversarial interactions.

The AI concierge can manage deliveries, greet guests and deflect inappropriate requests, all while maintaining a courteous and secure demeanor. Rigorous testing ensured the system's resilience against privacy threats and unanticipated scenarios. This customizable framework allows Swann to tailor settings to household-specific needs, enhancing user control.

Showcased at CES and named a Smart Home Honoree, this innovation sets the stage for Swann's future AI-powered products, demonstrating that cutting-edge security solutions can protect both homes and privacy in an increasingly connected world.

Actionable advice

Things to do (Adopt)

- **Use the Responsible Tech Playbook** for practical activities software delivery teams can actually do on projects.
- **Implement holistic and consistent policies around AI use.** Shadow AI without oversight can lead to a diverse range of issues, risking everything from privacy breaches to reputational problems.
- **Leverage new techniques to make generative AI more reliable.** This includes things like evals (a form of testing whereby outputs are assessed according to the context in which they will be received) and guardrails, a programmed set of policies that permits and prevents certain kinds of outputs. These need to be underpinned by your values and to do that you need to be able to define and articulate them.
- **Ensure your technology strategy is collaborative, not top down.** Although senior decision makers have an important role to play in setting a vision and an agenda for how an organization will use technology, involving other parties in that process not only builds organizational trust but also helps increase confidence in the decisions that are made. One effective way of doing this is to create a technology radar like the one we create twice a year at Thoughtworks. It allows people to question, voice concerns and propose alternatives in a way that is safe and supportive.
- **Treat responsible tech practices as a capability and skill issue.** Identify relevant training opportunities and resources for technology teams and other parts of the organization and ensure that teams see this as a valued area in which to develop new skills.
- **Be intentional about the social platforms you're using.** Are they spaces in which you want to play? Do you want to be associated with the type of content that is shared there?






Actionable advice

Things to consider (Analyze)

- **Dedicate time to revisiting and reconsidering your values.** Are they meaningful? Can they be practiced? Do they actually guide action?
- **Consider whether your organization is putting its professed values into action.** If it's not, why not? What are the business risks holding you back?
- **Think about how you might measure responsibility.** This could include anything from environmental measures through to employee perception and morale.
- **Employee attitudes to your values.** Do people feel they are being embraced? Are they something they themselves can enact?
- **Legislation.** Monitor new regulations and analyze how they may impact your organization. Keeping a close eye on broader conversations about future changes to law can also help you be prepared for the future and avoid being reactive.
- **Ownership and accountability inside the organization.** Many issues around responsibility quite literally require oversight from a responsible person. Expecting it will take care of itself is likely to be ineffective. Think seriously about who should be responsible or accountable and who should own or measure your performance in this area.
- **Pay attention to your software supply chain.** Do you understand what's in your stack? Consider using a software bill of materials (SBOM) to track dependencies and provide technological transparency.

Things to watch for (Anticipate)

- **How the law evolves on legal culpability for negative impacts of tech.** This is a question largely pertinent to AI, but it is also important to monitor legislation around content, data privacy and accessibility.
- **Consumer attitudes to AI.** AI is currently extremely hyped. However, this doesn't mean that consumer sentiment will follow the industry in the medium or long-term. Poor products, negative and dangerous effects and even pure fatigue may cause many people to see AI as either problematic or overhyped.
- **The future of ESG investing.** Just because ESG falls out of favor does not mean responsibility — whether that's social, environmental or otherwise — no longer matters. However, it may make it harder to make the case for it. Leaders need to do the right thing and ensure that there is a commitment to behaving with integrity and transparency — this will strengthen consumer confidence in your brand in the long-term.



Enabling richer experiences through multimodal interactions

Modes of interaction between people and machines have proliferated rapidly in recent years, to encompass text, voice, images, video, gestures — and with affective computing, even emotional cues — that may be unconscious.

This presents new possibilities, but also a challenge in requiring organizations to think through how different forms of interaction can be combined and used across the entire customer experience. Doing that can raise multiple questions. When is gesture or image recognition more helpful than a standard text interaction? When are physical buttons important, and in what situations could they constitute an obstacle? Carefully considering context and cognitive load can help ensure a positive experience that does not jar, irritate or unintentionally make an interaction more difficult. Innovation should never degrade a product or service.

While enterprises focus on enhancing their customers' online experiences they can gather insights that are applicable in other areas of the customer experience.

For instance, fashion retailers are increasingly using virtual fitting rooms to give shoppers a better sense of how their clothes will look when worn — by capturing their mix and match choices online, in-store staff can readily locate those items when the consumers come into the shop.



AI is playing a significant and expanding role not only in how we interact with services and products — most often through chatbots — but also how we model interactions. Agentic AI, for instance, can simulate human behavior to help sales teams identify prospects, or anticipate the customer response to a product or campaign.

Businesses should also remember not every new interaction is embraced by the customer. The backlash provoked by trends like gamification and addictive tech mean organizations need to be more considered in who they're building for, and why. This involves paying close attention to the friction in the customer experience and the accessibility implications of different channels for interaction, but also rethinking the foundations of what it means to interact with a product or brand. The best outcomes stem from avoiding the temptation to prioritize only optimization or stickiness, and considering more holistically how innovation in interaction can add value, or expand the effective range of a service.

“Human computer interaction is becoming more diverse than ever, which means organizations — in a huge range of industries — have the opportunity to create exceptionally rich experiences for everyone from their customers to their own employees. However, what is critical is ensuring experiences are always context specific, so thinking about what mode is most relevant when is vital.”


Chloe Blanchard
Principal Design Researcher, Thoughtworks




Signals

- **AI-driven chatbots becoming increasingly commonplace in domains such as customer service.** Research suggests the global chatbot market, valued at just over US\$5 Billion in 2023, will surge to over US\$36 Billion by 2032. Chatbots also continue to evolve, with techniques like retrieval-augmented generation (RAG) improving the accuracy and reliability of interactions.
- **Steady growth in the voice interaction space**, with more companies adopting voice assistants, and voice searches and requests accounting for a growing proportion of consumer activity.
- **Legislation taking aim at ‘dark patterns,’** aspects of the user experience that are subtly designed to nudge users into performing certain actions that are typically not to their benefit. [Article 25 of the European Union’s Digital Services Act](#) is one example. Similarly, accelerating and intensified gamification, as seen in the prevalence of controversial loot boxes, is coming under more government scrutiny.
- **Consumers growing more conscious of ‘screen time’ and actively seeking to avoid addictive tech.** There’s been a noticeable rise in apps to help users manage how they interact with their devices, or cut down interactions altogether, with Google recently rolling out ‘screentime reminders’ as part of the Digital Wellbeing suite for the Android mobile operating system. Australia banning under 16s from social media is one high-profile example of the rising fear around excessive screen exposure.

- **Agentic AI is being presented as the 'next frontier of generative AI.'** Able to handle more complex tasks, accept directions in natural language and work with other software tools, agents could change how we collaborate, complete tasks and interact with other technologies.
- **Continued innovation in AR and VR.** Despite the lukewarm response to some of the high-profile headsets that have launched, we still see companies making big bets in this space. AI integration is coming to Ray-Ban Meta smart glasses and Meta has also unveiled plans to release a neural wristband in the future, notably as part of its Orion project. Even if consumers have yet to embrace this technology fully, there's no shortage of companies seeking to find a killer app.
- **A neuralink chip being inserted into a person for the first time,** marking a major step forward in the development of direct brain-computer interfaces. The chip was also recently greenlit for clinical trials outside the US, as a potential means to enable paralyzed individuals to use the mind to engage with digital devices.
- **Consumer appetite for AR/VR retail experiences,** which iconic brands like Louis Vuitton have leveraged to engage customers and enhance their brands. One recent survey showed nearly a third of US consumers are keen for VR to recreate brick-and-mortar shopping routines.



Trends to watch



Seeing now

Adopt

1. Accessibility in multimodal experiences
2. AI agents
3. AI-assisted software development
4. Biometric authentication
5. Developer experience platforms
6. Digital twin
7. Green computing
8. Model training optimization
9. Privacy first

Analyze

10. Augmented reality
11. Automated workforce
12. Autonomous robots
13. Autonomous vehicles
14. Changing perceptions of AI
15. Consumer XR

16. Internet regulation
17. Mindful screen interaction
18. Multimodal interactions
19. Personalized healthcare
20. Satellite networks
21. Semantic representational technologies
22. Tactile interaction
23. Touchless interactions

Anticipate

24. Addictive tech
25. Understandable consent

Beginning to see

Adopt

26. Industrial XR

Analyze

27. GenAI computer control
28. Talk to data

Anticipate

29. Affective (emotional) computing
30. Brain-computer interfaces
31. Next-generation wearables

On the horizon

Adopt

—

Analyze

—

Anticipate

32. AGI research

The opportunities

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



Explore multimodal forms of interaction to create customer engagement breakthroughs.

Choices on which channels of interaction to adopt or offer need to be guided by context, situation, geography and user preferences. But examples like Canon's World Unseen exhibition, which featured images accessible to people with visual impairments, show how new technologies can expand what it means to interact with a product, service, or company, and deepen relationships with users or consumers.



Expand their potential market by bringing accessibility to bear.

Innovations like Signapse's automated sign language translation, and AI models that allow visually impaired people to interact with applications through natural language, show the potential for companies to open access to their brand or services for new customer groups.



Leverage AI agents to offer new conveniences to consumers.

Though agentic AI remains a nascent category, the technology is developing quickly, and can mimic human responses and behavior with uncanny levels of accuracy. This opens the possibility of automating more complex tasks and processes to help consumers accomplish their goals, whether by retrieving information, providing tailored advice or offering multilingual support.



Adopt extended reality (XR) to improve training and onboarding.

Organizations such as Boston's Mass General Brigham hospital are applying AR, VR and mixed reality to replicate real-life environments and situations, helping employees grow familiar with their roles and master new skills in a more immersive way. XR can be particularly valuable for training workers in high-risk industries, by exposing them to emergencies without any real-world danger.



Reimagine how teams, and customers, interact with data.

By applying AR and VR, organizations can move data visualization into new dimensions, providing users richer, more in-depth experiences of information, lightening the cognitive load and even potentially accelerating time to insights.



What we've done

Abraham Lincoln Presidential Library and Museum

Thoughtworks partnered with Google Cloud to transform the visitor experience at this important institution focusing on the life and legacy of Abraham Lincoln. AI, XR and VR technologies are being applied to make exhibits more immersive, by providing photorealistic experiences and access to additional layers of information. The project is also developing assisted visual guides to enhance accessibility for people with disabilities, and multilingual content for non-English speaking visitors.

Actionable advice

Things to do (Adopt)

- **Develop interactions with accessibility in mind**, and consider how new forms of interaction can support the creation of experiences that cater to all kinds of users. This may require educating and upskilling design and development teams, so they're aware of the expanding range of technologies that can be leveraged — or in some cases, rethinking how to approach the interaction.
- **Examine the impact of voice search on the way you manage and index products and information.** As technologies like GenAI-enabled agents make more queries and interactions verbal, the way organizations identify and categorize information will need to change. Optimization of data for voice search, and natural language metadata and tagging are likely to become increasingly important features of major platforms for interaction such as websites.
- **Prioritize quality, reliability and stability when developing new interactions**, and subject them to a comprehensive testing strategy. More sophisticated interactions can throw up new challenges, and even degrade a service or brand. Chatbots have been known to go off script and should be subjected to regular evaluations, especially when techniques like RAG are involved.




Things to consider (Analyze)

- **Choose the right interaction for the context.** It's easy to get excited about new technologies, but you should always be questioning whether they're the best fit for a particular use case, or how they'll improve the experience of end users. Sometimes gesture recognition isn't needed, and a simple button will do. As companies like carmaker Hyundai have found out, consumers often prefer 'old-fashioned' means of interaction as their inherent frictions can impart a sense of confidence, even joy.
- **Think through user mental models when using technologies like AI.** When rolling out something like a chatbot, it's vital to map out what people are likely to be expecting from an interaction, and what assumptions they'll bring to the table. Steps to help end-users write better prompts, or to explain outputs, can improve the overall user experience with new technologies.
- **Don't pursue stickiness at all costs.** The drive to optimize all experiences or create something addictive isn't just dangerous for users; it's often unsustainable and can provoke backlash. Rather than contributing to 'dark patterns,' consider the flipside: are there any opportunities to nudge users towards better, more cost effective, environmentally friendly or healthy decisions?

Things to watch for (Anticipate)

- **Evolving consumer tech opening new interaction platforms.** As products like Meta's Orion exit the experimental phase and enter the mainstream, companies need to be prepared to build out new forms of interaction that cater to their expanding user bases.
- **Tighter regulation** taking aim at dark patterns and other forms of interaction perceived as manipulative. Rather than simply rules to follow, such developments should be seen as an opportunity to build better interactions that are sustainable over the long term, foster trust with user bases and create commercial advantage over optimization-obsessed competitors.



Unlocking greater value from physical-digital convergence

The convergence of the physical and the digital is advancing in both industry and consumer tech — even if progress isn't always smooth, or readily evident.

One of the bigger recent developments in the space was the launch of Apple's Vision Pro in early 2024. While it garnered much fanfare initially, the excitement quickly dwindled, and adoption has proven disappointing. But that shouldn't be viewed as emblematic of the trend as a whole, which continues to develop in lower-profile ways that will present opportunities and gain traction beyond the current core base of power users.

One case in point is consumer health technologies, where a recent deal has valued rising star Oura, a maker of smart rings, at over US\$5 Billion. The popularity of Oura's elegant body-monitoring devices demonstrates there's significant potential for convergence in specific domains like personalized healthcare.

In some cases, hybrid experiences are almost business-as-usual.

Take examples such as video doorbells, digital menus or ride pooling services. They might not seem earth shattering, but they point to how easily consumers can adapt to the convergence of physical and virtual services and products.



Taking advantage of these opportunities will depend on a high level of domain understanding, and selectively implanting cross-reality innovations in ways that deliver tangible benefits — rather than building products or systems without clearly defined business or consumer use cases. Enterprises will also need to be sensitive to changing consumer expectations and emerging regulations such as the EU Data Privacy Act and Texas Data Privacy and Security Act, which restrict what companies can do with the consumer data garnered by connected devices.

Rather than a barrier, regulatory developments should be seen as a chance for proactive businesses to think through how they leverage more prevalent hybrid technologies safely and effectively. The EU Act may enable more players to participate, potentially growing the ecosystem. And by ensuring data is accurate, reliable and carefully managed, and conducting robust testing to ensure system reliability and resilience, enterprises can demonstrate transparency and build trust with their customer base.

Meanwhile, the physical-digital trend will get another boost as GenAI systems turn from consuming the available written data on the internet to integrating data from smart devices and wearables. Expect significant innovation in this space.

“It’s important that we are thoughtful about when, where and how we use technology — be it a new product, feature, leveraging GenAI or introducing IoT. More and more I’m seeing users and legislators holding businesses accountable especially when it comes to the collection, use and management of data, something at the very core of physical-digital convergence.”


Mackenzie Dysart
Delivery Principal, Thoughtworks




Signals

- **The mixed track record of consumer VR/AR tech**, which despite multiple high-profile launches in 2024, from the aforementioned Vision Pro to the Meta Quest 3S, has yet to truly take off. This is due in part to doubts about their interfaces, and the purpose such devices serve. The return of tactile buttons is another sign some consumer VR/AR devices may be racing ahead of what the market wants, or is ready for.
- **Success stories arising from domain-specific, business use cases.** The expansion of Saga Robotics, which develops multi-functional robots capable of treating the diseases that strike some crops, and BMW’s promising testing of humanoid robots to reduce employees’ strain in its assembly lines, show how leveraging convergence in a targeted way can deliver positive outcomes.

- **The rise of generative AI enhancing supply chain management**, by automating processes and expanding the types and quantity of data used to predict product demand and disruptions. That said, this is a relatively new field, and enterprises should monitor the extent to which such initiatives can deliver ROI.
- **The continued emergence of high-fidelity digital twins**. Their full impact may have yet to be realized, but developments like NVIDIA's powerful Earth 2 platform, which breaks new ground for climate and weather forecasting, show how digital twins are solving problems by allowing data to be visualized and manipulated in a more immersive way.
- **Industrial automation and IoT** introducing new potential vectors for cyberattacks, as they are integrated into more essential business or production processes.
- **New legislation around data collection and privacy**, often specifically targeting wearables. Jurisdictions like California and Colorado are leading the charge with laws that extend privacy protection to data gathered by health implants and wearable devices. More such moves are likely to follow, urging enterprises active in this space to solidify their governance strategies and adopt policies on ethical data use.



Trends to watch



Seeing now

Adopt

1. Biometric authentication
2. Digital carbon management
3. Digital twin
4. Privacy first
5. Software-defined vehicles

Analyze

6. Augmented reality
7. Autonomous robots
8. Autonomous vehicles
9. Hardware security
10. Personalized healthcare
11. Satellite networks
12. Tactile interaction
13. Touchless interactions

Anticipate

—

Beginning to see

Adopt

14. Industrial XR

Analyze

—

Anticipate

15. Affective (emotional) computing
16. AI in robotics
17. Brain-computer interfaces
18. Next-generation robotics
19. Next-generation wearables

On the horizon

Adopt

—

Analyze

20. AI avatars

Anticipate

—

The opportunities

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



Improve interactions for customers and employees by thinking through the accessibility implications of hybrid technologies and embedded systems. In addition to making some experiences more intuitive and impactful — for example, by creating training simulations that allow employees to get ‘hands-on’ with processes in a risk-free way — physical-digital convergence is opening new possibilities for the differently abled. Financial services is one area in particular where digital onboarding and biometric authentication are improving access to services for people who face mobility issues.



Embrace emerging regulation to build trust with consumers. Research suggests consumers are likely to abandon companies they don’t trust with their data, and gravitate towards those with transparent data practices.




Leverage automation, embedded systems and other cross-reality technologies to improve operational resilience and efficiency, and deliver value for the business or end-users. The application of digital twins in supply chain management demonstrates how cutting-edge simulations and real-time data can predict risks and identify areas of a network or system that are ripe for optimization.



Empower employees by adopting automation and smart devices, to take over repetitive or even dangerous tasks and enable teams to concentrate on more strategic or creative work. Companies like Microsoft are also adopting a hybrid approach to bridge the engagement and productivity gaps between remote and on-site teams.



Develop a deeper understanding of your operations and your customers. The data collected and shared by smart devices and embedded systems can unlock new levels of information and insight, to guide both tactical and strategic decision making. New data sources can also enable leaps forward in the organization’s use of AI and analytics — provided data is harnessed and used in a secure and sensitive way.



What we've done

Reece

Who wouldn't want their dream bathroom? But where to start? Research by Reece, Australia's leading supplier of bathroom products, revealed that over a third of renovators struggled to visualize the end result. This insight sparked innovation. Reece partnered with Thoughtworks to create a 3D bathroom planner, Imagin3D™, empowering customers to select products, visualize their new space and confidently make decisions. Showroom consultants and trades also use the tool to finalize contracts and proposals. In the first three months post-launch, over 30,000 customers turned their dream bathrooms into reality with Imagin3D™.

Actionable advice



Things to do (Adopt)

- **Treat data as the key enabler of physical-digital convergence.** Smarter devices adopted by consumers or in the work environment can vastly increase the amount of data available to the enterprise; however, it's only when this data is accurate, reliable and carefully governed that it can be used to inform decisions or credibly reconstruct real-world conditions. Embrace data engineering principles to ensure data travels where it is needed and can be leveraged with confidence.
- **Make domain knowledge the starting point for hybrid innovations and experiments.** Rather than attempting to apply new physical-digital solutions enterprise-wide, explore how they can improve or supplement parts of a specific business process. Applying these technologies in a limited way is more likely to yield results, and can also build the knowledge and momentum needed to support broader adoption.
- **Embed security and resilience in systems that span the physical and digital** by adhering to proven engineering practices, such as robust testing and continuous delivery. Having team structures that encompass embedded and backend development can smooth the path to production.



Things to consider (Analyze)

- **Consider how existing and new legislation might impact how your business uses data**, particularly data connected to individual consumers. Conduct an honest assessment of not just existing risks, but the risks that might emerge if regulations were to tighten further. Look for opportunities to improve trust and transparency.
- **Understand that automation, robotics and IoT are long-term infrastructure projects.** Making the most of convergence in complex environments such as manufacturing facilities will often require significant investment, the retooling of multiple roles, and skills that the organization may lack. It's important to consider these barriers and whether the organization has the capacity to deliver. The decisions leaders make now could have implications for the next five to 10 years.
- **Examine opportunities to improve experience** through convergence of the physical and digital — not just for customers or employees facing physical or location-related barriers, but for everyone. The use of biometrics to facilitate payments, and the adoption of contactless technologies to create a safer, less strain-inducing environment for workers, are good examples of convergence that can benefit a wide range of end users.

Things to watch for (Anticipate)

- **Tighter legislation around data privacy, and a hardening of consumer attitudes**, which may limit what the organization can do with data or its ability to pursue convergence-related opportunities. This is especially when it comes to sensitive areas like health, with research from Deloitte showing consumers are increasingly concerned about privacy breaches or excessive tracking linked to their wearable devices.
- **The evolution of the AR/VR market.** While these technologies remain outside the mainstream, it may be only a matter of time before a device emerges that kicks adoption into high gear, and creates viable commercial applications or use cases in the process.



Glossary

A

Accessibility in multimodal experiences:

The expansion of interactions beyond traditional interfaces to include XR, voice, image and gesture recognition, among others, brings new challenges in accessibility. Ensuring inclusivity in these contexts requires innovative design and testing to accommodate diverse user needs.

Addictive tech: Some applications are specifically designed to be addictive through the use of techniques like gamification and dark patterns. This is driven by fierce competition for eyeballs and engagement — and while there may be commercial reasons to adopt such an approach, an increasing awareness of the societal and environmental harms of addictive tech makes addictive tech a key issue organizations need to think seriously about.

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 simulations can help us understand complex phenomena such as the spread of diseases or protein folding.

AGI research: The concept of artificial general intelligence (AGI) refers to an AI system that possesses a broad range of capabilities across a range of intellectual tasks — it's often compared to human-level intelligence. Debates about the threshold for AGI remain, and research into ways of achieving it continues and will play a part in wider discussions about AI and humanity.

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. One of the most interesting manifestations of this trend are 'agentic assistants' in which AI agents are used to accomplish particular tasks in certain domains, like HR or CRM.

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-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.

Glossary

AI avatars: A digital representation of a person. The use of artificial intelligence allows the avatar to mimic the person it represents, thus making it ostensibly more convincing and lifelike.

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

AI in robotics: Bringing today's AI capabilities into robotics is creating new levels of 'intelligence'. It can help robots better respond to situations and external stimuli and ostensibly make decisions about what actions to take in relation to its environment.

AI in security: AI is today often 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 observability: AI systems are notoriously opaque. Their complexity can make it very difficult to determine the relationship between inputs and outputs. AI observability is the broad practice of monitoring and analyzing an AI system's behaviors and performance to increase understanding and confidence that it is working as intended.

AI-ready data: AI-ready data is data that has been structured and organized in a way that makes it easy for it to be integrated with AI systems. It has a number of specific qualities: high-quality (auditable and verifiable), consistent across different platforms and robust, comprehensive metadata.

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 and is beginning to be extended to consumer interactions with AI.

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 (AR): 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, automation simplifies reporting by sifting through data. Increasingly, though, AI is 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.

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

Biometric authentication: A way of verifying an individual's identity that uses fingerprint, facial recognition or other similar technologies. It is today a valuable cybersecurity tool in many different domains and industries.

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

Changing perceptions of AI: AI technologies have been widely hyped and are therefore extremely visible not just in the industry but in wider society and culture. This means that attitudes and understandings of it — whether that's enthusiasm and excitement or distrust — are necessarily important to organizations that decide to use it. The pace of technological change, moreover, means that attitudes could also change quickly.

Collaboration ecosystems: When individuals or organizations share common goals, they will likely want to work together. To do so, though, 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. This is what a collaboration ecosystem is: it allows people to solve problems together.

Consumer XR: Consumer XR refers to products and services that give users extended reality experiences. High-profile devices like the Apple Vision Pro are shaping consumer XR, but the field is highly dependent on innovations in retinal resolution to ensure properly immersive experiences.

Context-aware systems: Systems that dynamically adapt their behavior using real-time contextual information, such as user location, activity, or preferences. While the concept has existed since the early days of ubiquitous computing, advancements in AI, IoT, and edge computing have significantly enhanced their capabilities. Modern context-aware systems deliver highly personalized and responsive experiences, becoming a competitive advantage across industries and signaling their importance towards adaptive and human-centric technologies.

D

Data catalog: A comprehensive inventory of an organization's data assets. Crucially, it is built on well-organized metadata, which makes it easier for organizations to discover and retrieve a particular asset and then use it appropriately.

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 fitness functions: Automated tests that assess the quality, consistency and reliability of data in real time. By continuously assessing key characteristics, these functions ensure data meets predefined governance standards and remains fit for use in evolving workflows, facilitating interoperability and trust across data systems.

Data lineage: 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.

Glossary

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 architectures: Use of multiple data stores instead of singular, monolithic centralized stores. A good example is data mesh.

Decentralized personal data stores: A data architecture style where individuals control their own data in a decentralized manner, allowing access on a per-usage basis (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.

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

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.

Digital carbon management: Measuring organizational greenhouse 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 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.

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. The benefits of edge computing include reduced latency for real-time systems and improved data privacy. It's also possible to run AI/ML models at the edge too.

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.

Evaluating and managing AI outputs: Ensuring the quality, reliability, and safety of AI-generated outputs through evaluation frameworks — 'evals' — and guardrails. These include systematic tests to measure performance and tools to enforce ethical and operational standards, helping businesses deploy AI responsibly and effectively.

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

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 computer control: A new capability of generative AI tools to execute and automate computer-based tasks through natural language. They enhance digital workflows by enabling intuitive, conversational interactions with operating systems and applications. Examples include Claude's "computer use" feature and Auto-GPT, among others.

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

Green computing: Green computing is a diverse collection of practices and techniques that attempt to address the environmental impact of computation. It includes green cloud, green UX and green software development, all of which optimize systems, code and other part of technology infrastructure to improve computational efficiency and reduce waste.

H

Hardware security: The growth in smart devices and embedded systems have made hardware an even bigger target for cybercriminals and malicious actors. Ensuring hardware is secure is today a key step in ensuring security across the enterprise.

I

Impact funds: Impact funds, or impact investing, is a trend whereby investors target businesses tackling significant social or environmental challenges in a bid to both develop a solution and, in doing so, unlock substantial financial returns.

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.

Integrating unstructured data: Set of techniques and tools for processing and incorporating unstructured data, such as text, images, and videos, into workflows and decision-making. Approaches like natural language processing, computer vision, and data indexing systems make this data more accessible and actionable for businesses.

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

Interfacing with AI: Establishing standardized methods for integrating generative AI into business systems using tools like LLM proxies and OpenAPI. LLM proxies act as intermediaries that simplify AI interactions, while OpenAPI defines clear, consistent interfaces for connecting AI models to applications, ensuring scalability and ease of use.

Internet regulation: The regulation of the internet has become more and more significant in recent years. This manifests itself in many different ways, from attempts to address harmful content, restricting children's use of social media and rules about how consumer data can be collected and used.

K

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

L

LLMOps: The practice of integrating LLMs into business operations, focusing on deployment, monitoring, security, and governance. This includes tools and processes for fine-tuning, performance tracking, cost management, and ensuring responsible AI use.

M

Mindful screen interaction: A growing shift toward intentional and balanced device use, driven by increased awareness of screen time. Tools like screen time trackers and focus apps exemplify this trend, supporting users in managing their digital habits.

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.

Model training optimization: Strategies and techniques to enhance the efficiency and effectiveness of machine learning model training. Examples include retrieval-augmented generation (RAG), which combines data retrieval with generative AI for precise outputs; causal inference, which identifies cause-and-effect relationships to improve generalizability and reduce training data requirements; transfer learning, which leverages pretrained models for faster adaptation; and automated hyperparameter tuning, which optimizes model performance with minimal manual effort. These approaches are crucial for reducing costs, minimizing energy consumption, and accelerating deployment.

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.

Multimodal interactions: Systems that enable users to interact through multiple input methods, such as text, voice, image and gesture recognition. By combining these modalities, tools and applications create more intuitive and accessible experiences across diverse contexts.

N

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.

Next-generation robotics: The next generation of robotics is underpinned by advancements in artificial intelligence and machine learning. These technologies are helping to bring new dimensions of responsiveness and 'reasoning' to robotics.

Next-generation wearables: The next generation of wearables are getting smaller but also ostensibly smarter thanks to the increasing integration of AI. These devices — ranging from the popular Oura to the Humane AI pin — offer users new ways to quantify the self.

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.

P

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-enhancing technologies (PETs): A collection of technologies and techniques designed to preserve user privacy while enabling secure and trustworthy interactions. Examples include anonymization, encrypted computing, differential privacy, decentralized identity (DiD) for self-owned digital IDs and verifiable credentials, and zero-knowledge proofs, which allow validation without exposing sensitive data. These tools play a critical role in safeguarding privacy in increasingly data-driven and interconnected systems.

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: The use of probabilistic states of photons, rather than binary ones and zeros, to execute algorithms with significant speedup in specific problem domains. Recent advancements, such as Google's breakthroughs in quantum error correction, signal progress toward scalable systems. However, these developments also raise concerns about security, as quantum computers could potentially break classical cryptographic protocols, driving interest in quantum-resistant encryption methods.

R

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.

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.

Semantic representational technologies:

A collection of techniques aimed at helping machines better understand data. It aims to put meaning at the very center of data, so concepts, categories and relationships can be better ‘understood’ by machines. For users, this can make it easier to search and manage incredibly complex data sets.

Small language models: An alternative to large language models (LLMs) that are more lightweight and efficient. While they aren’t as powerful compared to their larger siblings, because they require less memory and computational power they can be used in devices at the edge of a network.

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.

Synthetic data: Artificial data that mimics ‘real’ data. It is created algorithmically, expanding the potential size of a data set without requiring further data collection. This has many applications, from drug research to testing, and also has the benefit of reducing the risks and challenges that come from acquiring new, ‘real’ data.

T

Tactile interaction: Tactile interaction is an emerging trend in extended reality. It uses something called haptic feedback to enable richer and more immersive experiences where users can physically experience a virtual environment.

Talk to data: Talk to data (T2D) is a technology that allows users to interact with and analyze data using natural language queries as opposed to, say, the kinds of analytics and business intelligence dashboards that have become commonplace over the last two decades. It makes it easier to uncover insights and has a lower barrier to entry, giving more employees the ability to explore and ask questions about data.

Touchless interactions: The ability to interact with devices without touching. Specific technologies include hand tracking and voice and gesture recognition.

U

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.

Thoughtworks is a global technology consultancy that integrates strategy, design and engineering to drive digital innovation. We are over 10,000 Thoughtworkers strong across 48 offices in 19 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.