


Deloitte.
Insights

Tech Trends 2023





Trend 2 **Opening up to AI:** **Learning to trust our AI colleagues**

While the value of artificial intelligence is now undoubtable, the question has become how to best use it—and that often boils down to how much workers and end users trust AI tools.

Computers were once seen as more or less infallible machines that simply processed discrete inputs into discrete outputs, whose calculations were never wrong. If a problem ever arose in a calculation or business process, it was definitionally caused by human error, not the computer.

But as machines encroach on ever-more humanlike tasks that go beyond basic number crunching and enter the realm of discernment and decision-making via artificial intelligence (AI), the business world is developing a new understanding of what it means to trust machines.

The degree to which businesses and workers learn to trust their AI “colleagues” could play an important role in their business success. Most organizations today say they’re data-driven. Many even call themselves AI-fueled companies.¹ There’s plenty of evidence suggesting businesses that use AI pervasively throughout their operations perform at a higher level than those that don’t: Enterprises that have an AI strategy are 1.7 times more likely to achieve their goals than those that lack such a vision.²

Yet the underlying AI tool implemented in a given workflow matters less.³ With cloud vendors increasingly offering prebuilt models, any business can access world-class AI functionality with

a few clicks. The top-performing facial recognition vendors ranked by the National Institute of Standards and Technology deliver comparable performance, and they’re all easily accessed through cloud-based services.⁴ It’s what you do with the tool that’s important—and whether your people, customers, and business trust the results.

So what may matter in the future is not who can craft the best algorithm, but rather who can use AI most effectively. As algorithms increasingly shoulder *probabilistic* tasks such as object detection, speech recognition, and image and text generation, the real impact of AI applications may depend on how much their human colleagues understand and agree with what they’re doing.

People don't embrace what they don't understand. We spent the last 10 years trying to get machines to understand us better. Now it looks like the next 10 years might be more about innovations that help us understand machines.

Developing processes that leverage AI in transparent and explainable ways will be key to spurring adoption.

"What we're designing is an interface of trust between a human and a machine," says Jason Lim, identity management capability manager at the [Transportation Security Administration](#). "Now you're taking an input from a machine and feeding it into your decision-making. If humans don't trust machines or think they're making the right call, it won't be used."⁵

Think of deploying AI like onboarding a new team member. We know generally what makes for effective teams: openness, rapport, the abil-

ity to have honest discussions, and a willingness to accept feedback to improve performance. Implementing AI with this framework in mind may help the team view AI as a trusted copilot rather than a brilliant but taciturn critic. When applications are transparent, resilient, and dependable, they can become a natural part of the workstream.

Now Business-critical but inscrutable

When recruiting new team members, managers often look for the right mix of skills and fit. Few leaders doubt AI's abilities to contribute to the team. According to one survey, 73% of businesses say AI is critical to their success.⁶

But they're less sold on fit. Currently, enterprises have a hard time trusting AI with mission-critical tasks. The same report found that 41% of technologists are concerned about the ethics of the AI

tools their company uses, and 47% of business leaders have concerns about transparency,⁷ the ability for users to understand the data that went into a model.

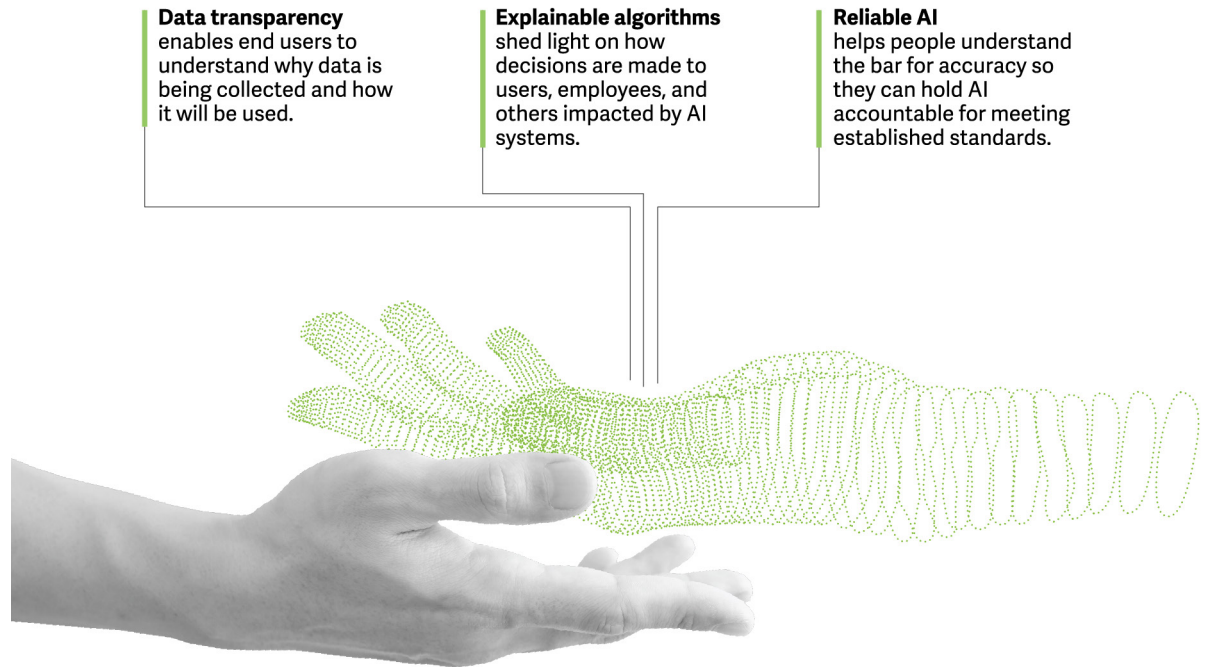
Enterprises are also grappling with a related concept, explainability, the ability of a model to give an explicit justification for its decision or recommendation. Explainability in AI systems is necessary when it is required by regulations, but it's also becoming expected functionality in situations where it helps make clear to end users how to use a tool, improve the system generally, and assess fairness.⁸ Explainability is one of the biggest differentiators between the successful use of AI at scale and failure to reap returns on AI investment, yet many businesses haven't figured out how to achieve it.

New

From black box to glass box

Mistrust of AI can come from business leaders, front-line workers, and consumers. Regardless of its origin, it can dampen enterprises' AI enthusiasm and, in turn, adoption. But leading organizations are working on solving issues that diminish trust in AI implementations. Some of the most effective approaches treat AI not so much as a point technology but rather as a piece in a larger process, considering the various stages where humans interact with the AI system and working to identify and address areas of potential mistrust. Acknowledging that AI tools are techniques to be woven into the larger tapestry of processes within an organization can make it easier to fix trust issues proactively. For more trusted AI, forward-thinking enterprises are leaning on data transparency, algorithmic explainability, and AI reliability (figure 1).

FIGURE 1: How to make AI more trusted



Source: Deloitte analysis.

Data transparency

Transparent data-collection methods enable the end user to understand why certain pieces of information are being collected and how they're going to be used. When users have this control, they can make informed decisions about whether the AI tool represents a fair value exchange.⁹

The Saudi Tourism Authority used this approach when developing a new application for travelers. The app uses AI to guide tourists through their stay in the country, recommending restaurants, attractions, and other activities based on location and preferences. But importantly, the user is in control of the data they provide to the app. Visitors can determine how much or how little data they hand over, or can opt out completely, with the understanding that giving the app less data access may mean less-tailored recommendations.¹⁰ This stands in contrast to many apps that have all-or-nothing data access requirements that generally serve as a poor foundation for trust.¹¹

Algorithmic explainability

One of the biggest clouds hanging over AI today is its black-box problem. Because of how certain algorithms train, it can be very difficult, if not impossible, to understand how they arrive at a recommendation. Asking workers to do something simply because the great and powerful algorithm behind the curtain says to is likely to lead to low levels of buy-in.



One automaker in the United Kingdom is tackling this problem by bringing front-line workers into the process of developing AI tools. The manufacturer wanted to bring more AI into the vehicle-assembly process by enabling machine learning to control assembly robots and identify potentially misaligned parts before the vehicle gets too far into the assembly process. At the start of the development process, engineers bring in frontline assembly workers to gauge their perception of problems and use that to inform development. Rather than dropping AI into an arbitrary point in the production process, they use it where the assemblers say they most need help.

The tools ultimately built are interpretable because the workers' input forms the basis of alerts and recommendations. In other words, it's easy for assemblers to see how the AI platform's recommendations map to the problems they themselves helped define. By bringing in

workers at the start and helping them understand how the AI functions, developers are able to support the assembly team with trusted cobot coworkers rather than a silicon overlord dictating opaque instructions.

AI reliability

People have grown accustomed to a certain level of reliability from work applications. When you open an internet browser or word-processing application, it typically simply “behaves.” More specialized business applications such as customer relationship management platforms and enterprise resource management tools may be a bit more finicky, but their challenges are fairly well established, and good developers know how to troubleshoot them.

With AI, the question isn’t whether it will work but rather how accurate the result will be or how precisely the model will assess a situation. AI is generally neither right nor wrong in the traditional sense. AI outputs are probabilistic,

expressing the likelihood of certain outcomes or conditions as percentages—like a weather forecast predicting a 60% chance of rain—which can make assessing reliability a challenge. But workers need to know how accurate and precise AI is, particularly in critical scenarios such as health care applications.¹²

AI is sometimes viewed as much as an art as a science, but that may need to change for robust adoption. Organizations that take a rigorous approach to ensuring AI reliability consistently see better results. Those that document and enforce **MLOps processes**—a set of procedures designed to ensure machine learning tools are deployed in a consistent and reliable manner—are twice as likely as those that don’t to achieve their goals and to deploy AI in a trustworthy way.¹³ Taking an operations-minded approach puts guardrails around AI and helps build confidence that it is subject to the same standards of reliability as any other business application.

But reliable doesn’t necessarily mean perfect. Just as human coworkers will never deliver perfect results every time, AI too will make mistakes. So the bar for reliability is not perfection, but rather how often it meets or exceeds an existing performance standard.

We spent the last 10 years trying to get machines to understand us better. Now it looks like the next 10 years might be more about innovations that help us understand machines.

Next

Creative machines

As enterprises deploy AI in traditional operational systems, a new trend is taking shape on the horizon: generative AI. We're already seeing the emergence of tools such as OpenAI's Dall-E 2 image generator and GPT-3 text generator. There's a generative model for music called Jukebox that lets users automatically create songs that mimic specific artists' styles.¹⁴ AI is increasingly being used to automatically caption live audio and video.¹⁵ These types of content generators are getting more sophisticated by the day and are reaching the point where people have a hard time telling the difference between artificially rendered works and those created by humans.

Concern over automation's impact on jobs is nothing new, but it is growing ever more pronounced as we head toward this automatically generated future. In many cases, generative AI

is proving itself in areas that were once thought to be automation-proof: Even poets, painters, and priests are finding no job will be untouched by machines.

That does not mean, however, that these jobs are going away. Even the most sophisticated AI applications today can't match humans when it comes to purely creative tasks such as conceptualization, and we're still a long way off from AI tools that can unseat humans in jobs in these areas. A smart approach to bringing in new AI tools is to position them as assistants, not competitors.

Companies still need designers to develop concepts and choose the best output, even if designers aren't doing as much of the manipulating of images directly. They need writers to understand topics and connect them to readers' interests. In these cases, content generators are just another tool. As OpenAI's CEO Sam Altman writes in a blog on DALL-E-2, "It's an example of a

world in which good ideas are the limit for what we can do, not specific skills."¹⁶

Workers and companies that learn to team with AI and leverage the unique strengths of both AI and humans may find that we're all better together. Think about the creative, connective capabilities of the human mind combined with AI's talent for production work. We're seeing this approach come to life in the emerging role of the prompt engineer.¹⁷ This teaming approach may lead to better job security for workers and better employee experience for businesses.

AI continues to push into new use cases through emerging capabilities that most people thought would remain the exclusive domain of humans. As enterprises consider adopting these capabilities, they could benefit from thinking about how users will interact with them and how that will impact trust. For some businesses, the functionality offered by emerging AI tools could be game-changing. But a lack of trust could ultimately derail these ambitions.

Endnotes

1. Beena Ammanath et al., [Becoming an AI-fueled organization: State of AI in the enterprise, 4th edition](#), Deloitte Insights, October 21, 2021.
2. Ibid.
3. Abdullah A. Abonamah, Muhammad Usman Tariq, and Samar Shilbayeh, [“On the commoditization of artificial intelligence,”](#) *Frontiers*, September 30, 2021.
4. Patrick Grother et al., [Face recognition vendor test \(FRVT\)](#), National Institute of Standards and Technology, July 2021.
5. Deloitte, [The Transportation Security Administration makes digital transformation human](#), Deloitte Insights, October 5, 2022.
6. Appen, [The state of AI and machine learning](#), accessed October 26, 2022.
7. Ibid.
8. Reid Blackman and Beena Ammanath, [“When — and why — you should explain how your AI works,”](#) *Harvard Business Review*, August 31, 2022.
9. Irfan Saif and Beena Ammanath, [“Trustworthy AI is a framework to help manage unique risk,”](#) *MIT Technology Review*, March 25, 2020.
10. Deloitte, [Saudi Arabia’s digital government stays ahead of the curve: How a nationwide technology innovation ecosystem is enhancing the digital government experience for citizens—and staying focused on the future](#), Deloitte Insights, October 28, 2022.
11. Catharine Bannister and Deborah Golden, [Ethical technology and trust: Applying your company’s values to technology, people, and processes](#), Deloitte Insights, January 15, 2020.
12. Saif and Ammanath, [“Trustworthy AI’ is a framework to help manage unique risk.”](#)
13. Ammanath et al., [Becoming an AI-fueled organization](#).
14. Prafulla Dhariwal et al., [Jukebox: A generative model for music](#), Cornell University, April 30, 2020.
15. IBM, [“Closed captioning software: Leverage AI with speech recognition for automatic captioning on live broadcasts and online video,”](#) accessed October 26, 2022.
16. Sam Altman (blog), [“DALL•E 2,”](#) April 6, 2022.
17. Tori Orr, [“So you want to be a prompt engineer: Critical careers of the future,”](#) *VentureBeat*, September 17, 2022.



Sign up for Deloitte Insights updates at www.deloitte.com/insights



Follow @DeloitteInsight

www.deloitte.com/us/TechTrends



Follow @DeloitteOnTech

Deloitte Insights contributors

Editorial: Aditi Rao, Blythe Hurlay, Rupesh Bhat, Aishwarya Iyer, and Emma Downey

Creative: Jim Slatton, Adrian Espinoza, Sylvia Chang, Alexis Werbeck, Molly Woodworth, Jaime Austin, Natalie Pfaff, and Heidi Morrow

Deployment: Pooja Boopathy

Cover artwork: Found Studio

About Deloitte Insights

Deloitte Insights publishes original articles, reports and periodicals that provide insights for businesses, the public sector and NGOs. Our goal is to draw upon research and experience from throughout our professional services organization, and that of coauthors in academia and business, to advance the conversation on a broad spectrum of topics of interest to executives and government leaders.

Deloitte Insights is an imprint of Deloitte Development LLC.

About this publication

This publication contains general information only, and none of Deloitte Touche Tohmatsu Limited, its member firms, or its and their affiliates are, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your finances or your business. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser. None of Deloitte Touche Tohmatsu Limited, its member firms, or its and their respective affiliates shall be responsible for any loss whatsoever sustained by any person who relies on this publication.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee ("DTTL"), its network of member firms, and their related entities. DTTL and each of its member firms are legally separate and independent entities. DTTL (also referred to as "Deloitte Global") does not provide services to clients. In the United States, Deloitte refers to one or more of the US member firms of DTTL, their related entities that operate using the "Deloitte" name in the United States and their respective affiliates. Certain services may not be available to attest clients under the rules and regulations of public accounting. Please see www.deloitte.com/about to learn more about our global network of member firms.

Copyright © 2023 Deloitte Development LLC. All rights reserved.

Member of Deloitte Touche Tohmatsu Limited