Deloitte.

Toward humanity's brightest future with Generative AI

Deloitte AI Institute™

About the
Deloitte AI Institute™

The Deloitte AI Institute helps organizations connect the different dimensions of a robust, highly dynamic and rapidly evolving AI ecosystem. The AI Institute leads conversations on applied AI innovation across industries, with cutting-edge insights, to promote human-machine collaboration in the "Age of With".

The Deloitte AI Institute aims to promote a dialogue and development of artificial intelligence, stimulate innovation, and examine challenges to AI implementation and ways to address them. The AI Institute collaborates with an ecosystem composed of academic research groups, start-ups, entrepreneurs, innovators, mature AI product leaders, and AI visionaries, to explore key areas of artificial intelligence including risks, policies, ethics, future of work and talent, and applied AI use cases. Combined with Deloitte's deep knowledge and experience in artificial intelligence applications, the Institute helps make sense of this complex ecosystem, and as a result, delivers impactful perspectives to help organizations succeed by making informed AI decisions.

No matter what stage of the AI journey you're in; whether you're a board member or a C-Suite leader driving strategy for your organization, or a hands on data scientist, bringing an AI strategy to life, the Deloitte AI institute can help you learn more about how enterprises across the world are leveraging AI for a competitive advantage. Visit us at the Deloitte AI Institute for a full body of our work, subscribe to our podcasts and newsletter, and join us at our meet ups and live events. Let's explore the future of AI together.
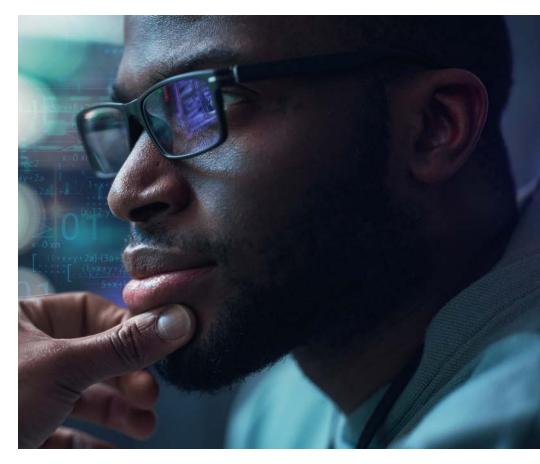
www.deloitte.com/us/AIInstitute

We are at a pivotal moment in the history of human invention. For years, the confluence of technologies that power the modern world, like analytics and artificial intelligence (AI), have been changing and challenging how the world works and societies function.

The emergence of Generative AI (GenAI) is a significant leap forward and one of the fastest-moving innovations in recent memory. Its potential is at the top of every business agenda, as are questions about how to approach and use this transformative technology.

Many of us are energized by the opportunity to be on the ground floor of shaping the future. We are already seeing GenAI use cases that drive process efficiency, improve automation, and enhance how humans and machines work together in the enterprise. Yet, these applications only scratch the surface of GenAI's potential. We can make tomorrow fantastically bright, but we should be actively weighing the societal implications of GenAI implementation, particularly as it relates to the future of work, equity and diversity in talent, and promoting trust in technology.

If our attention at this inflection point is narrowly focused on business benefit, we may miss the broader potential impact of GenAI on people and society at large. Purpose-driven organizations that aspire to balance business success with making a positive and lasting impact should take note. This is an opportunity—and an obligation—to lead, and it begins with an appreciation of the kind of technology we have in our hands and how it can transform the world for the better.

A grand vision for Generative AI and society

GenAI is acclaimed for its capacity to retrieve and assemble information in a way that appears to be human-created. Text, images, sounds, even 3D models can be conjured with a prompt and without much technical knowledge of how it functions.

To this point, much of the experimentation in using GenAI has focused on automation that reduces costs and improves efficiency. Use cases such as call centers, rapid content creation for marketing and engagement, and the ability to reduce manual and time-consuming tasks can deliver value for organizations.

These kinds of applications, however, are the low-hanging fruit of AI deployment, which have been targeted for years with more traditional AI models. **With GenAI, we are not simply contending with another business tool. The stunning potential is bigger than that, and with it, we can affect positive impact in people's lives for generations to come.**

Imagine an elderly relative who develops a life-threatening disease, an experience unfortunately common for many of us. Traditionally, doctors might conduct a variety of tests, use their expertise to recommend a treatment, and hope it works. If it does not, the next treatment is tried. In this pre-AI environment, finding the right health solution generally hinges on the physician's knowledge and trial and error.

What if health providers leveraged GenAI in the search for a life-saving solution? With it, doctors could more quickly and thoroughly explore vast databases of existing treatments and medical insights, helping to find options that may otherwise remain hidden in the data. Physicians could leverage GenAI to access all relevant medical studies and pharma trials, with the insights summarized in an easily digestible format. In some instances, GenAI could be used to model a pharmaceutical treatment precisely tailored to the patient's genome.

These kinds of applications can enable doctors to make better treatment decisions, faster. The short-term outcome is ideally an effective treatment that saves the patient's life, but picture this medical future at scale. We are at the doorstep of a world where every person can receive this level of informed, patient-centered treatment. When we consider the good GenAI can do, this is the level of global impact we should imagine.
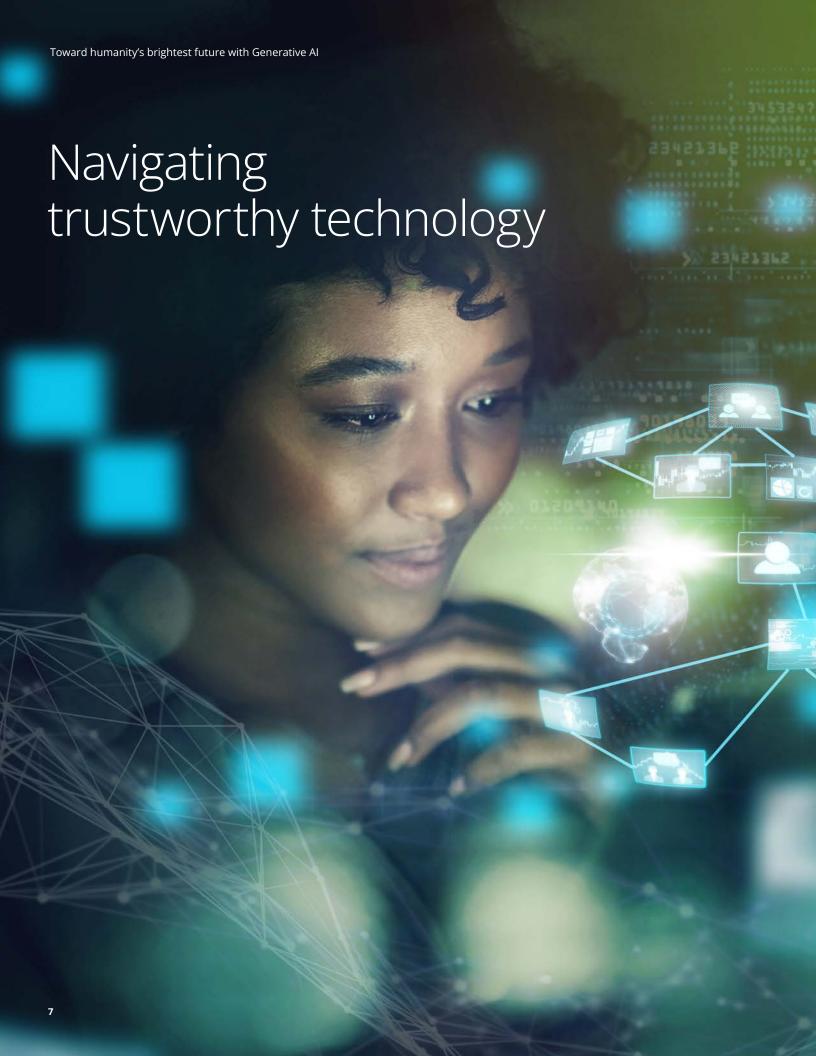
We can anticipate a similar scale of impact in access to education. In many places around the world, even attending primary education can be a challenge for millions of young people. There may be language barriers to exploring digitized information, much of which is written in English. There can also be practical hurdles, such as having to travel long distances to reach a school, the inability to afford it, and limited capacity in education systems. AI, and specifically GenAI, can minimize these hurdles and drive educational equity.

Imagine a young person in a rural area who lacks the opportunity to attend school. With only an internet connection, GenAI could help them access lessons and activities from teachers across the world, irrespective of language. They could translate educational content into their native dialect, with GenAI helping to condense and explain new material. We can even envision GenAI assistants who help educators expand capacity and reach more curious minds. When students have commensurate opportunities to learn, no matter who or where they are, the barriers to education diminish. **Do we not have an obligation to act in a way that combines the power of human collaboration with the power of technology to build that future?**

These are just two examples to demonstrate how we should be thinking of GenAI as a tool for the amplification of human potential. An essential catalyst to reach that potential is trust.

Navigating trustworthy technology

Risk management has been an enterprise priority for decades, and **in recent years, the volume and variety of technologies creating risk have only multiplied.**

Just as organizations and societies were getting their arms around the risks traditional AI can present, along came GenAI, revealing a vast landscape of new, even more profound concerns. It is no longer simply the risks of the technology itself that must be managed. Now, we are contending with implications that can permeate all elements of an organization and society more broadly.

There is no shortage of forecasts pointing out known and potential GenAI risks, from the threat of misinformation at scale to latent biases shaping decision making to even more existential problems. **We are also beginning to see how the challenge with GenAI is not erasing risk, but balancing the benefits against risks in a way that generates trust.**

Leaders of purpose-driven organizations already understand that their decisions directly impact society in the short- and longer-terms. What we do now has a direct impact on the future we, and more importantly our descendants, will inhabit. How we prepare ourselves, our organizations, and our societies to use GenAI in a responsible way will influence trust in the technology itself, and by extension, the brightest futures we hope to see come to fruition.

It starts with a human-centric vision. We can be tempted by the power of technology to drive our decision making, to turn over human activities to a machine and even outsource our creativity. This leads to a future where humans are not at the center of the vision, where our emotional intelligence is barely reflected and machine capabilities are misaligned with human values. In this, our humanity could be sidelined, and we could begin to lose sight of a GenAI future built around what makes us unique. **If we reject a machine-centric vision and maintain an allegiance to humanity first, we see that a priority for GenAI use is equity in value.**

If enterprise leaders pursue GenAI application in a way that only replaces humans and fills the creative void with image generators and chatbots, it will harm trust. Skepticism will grow. People in communities all over the world will rightly ask, is GenAI simply a tool for personal gain and business growth or is it for the good of society? That question, fueled by public suspicion over how organizations choose to use this technology, has the power to dim the brightest future.

One pathway to promoting equity is inclusiveness in GenAI development and usage. Take an enterprise example. An organization's data scientists are primarily focused on building, deploying, and improving GenAI use cases. Could a dozen or so technologists with little bandwidth possess the lived and professional experience to imagine all the ways GenAI could be used and the risks that may arise? No. Conversely, consider the volume and quality of ideas that could emerge from dozens of stakeholders from across the enterprise coming together for dynamic and creative discussions on GenAI value and risk. Human centricity is a natural byproduct of an inclusive and collaborative approach to conceiving and using GenAI.

By mobilizing more people with a greater diversity of experiences and perspectives, we can build toward a GenAI-fueled future where all people see their needs and concerns reflected in the most powerful technologies of the day. This human-centric GenAI vision moves toward amplifying (rather than dampening) the qualities and capabilities that make people special. A world of AI applications working in harmony with our emotional intelligence can advance our dreams, invigorate our passions, and stimulate our creativity. It can help us develop and grow as humans, and in turn, find meaning and purpose.

With the values of empathy, collaboration, and human communication as a foundation for GenAI, trust—and therefore, transformative and beneficial outcomes— will blossom.

To get there, we need business cultures that welcome diversity and grow by virtue of a range of backgrounds and perspectives. **We also need governance models that establish guardrails and responsibilities around GenAI deployment. Boards have an important role to play in this regard.** Board members are positioned to shape the longer-term vision for GenAI, with a focus on sustainable value and trustworthy and responsible application. They can, and should, ask probing questions about the impact on the organization and society and set strategic priorities to navigate those impacts. These activities inform model governance, including roles and responsibilities, post-deployment GenAI management, and legal and compliance matters.

Boards and executives working together can also probe how to ensure people, both in the workforce and elsewhere, have the technical capabilities and knowledge to keep pace with and use AI innovation. We need to support more people in understanding, using, and analyzing GenAI, optimizing the full potential of their skillsets such that they are equipped to join the conversation on how we create tomorrow for the benefit of all. Enterprises and their leaders have an evident role to play.

History is watching

Today, we look back on our predecessors and try to understand what their world looked like and **how and why they made the decisions they did.**

We look at the industrial era and try to fathom how child labor, unchecked pollution, and unsafe work environments were ever tolerated. We reflect on the mass production of the combustion engine and wonder why more people weren't concerned with emissions and the impact at scale on our planet. There are so many examples where, with hindsight, we struggle to grasp what our ancestors were thinking.

Our descendants will wonder the same about us. They will perhaps ask, did we make the best decisions in the early days of GenAI? Did we live in a world of greed, focused only on things that drive near-term benefits for ourselves, or were we motivated to make the world a better place? Were we living in fear of a technology that creates risk, and did we let that lead us to erratic choices? Did we place people at the heart of our AI efforts, with the intention of magnifying and promoting those things that make us uniquely human? There is a common thread of responsibility that ties us to the future we create. **We must balance the trade-offs between the value we can access and the value we think is important for society.**

As leaders, we set the tone. By infusing trust, diversity, and ethics in all our decision making, we build the stage for GenAI to be used for the benefit of all and in the interest of the brightest futures for the next generation.

It is up to us to get this right.

Reach out for a conversation.

Lara Abrash
Chair
Deloitte US
labrash@deloitte.com

Beena Ammanath
Global Deloitte AI Institute Leader
Deloitte LLP
bammanath@deloitte.com

Deloitte.

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.