



Ethical AI

Five guiding pillars



About the authors



Todd Lohr

Principal, Advisory,
KPMG Digital Lighthouse Network Leader

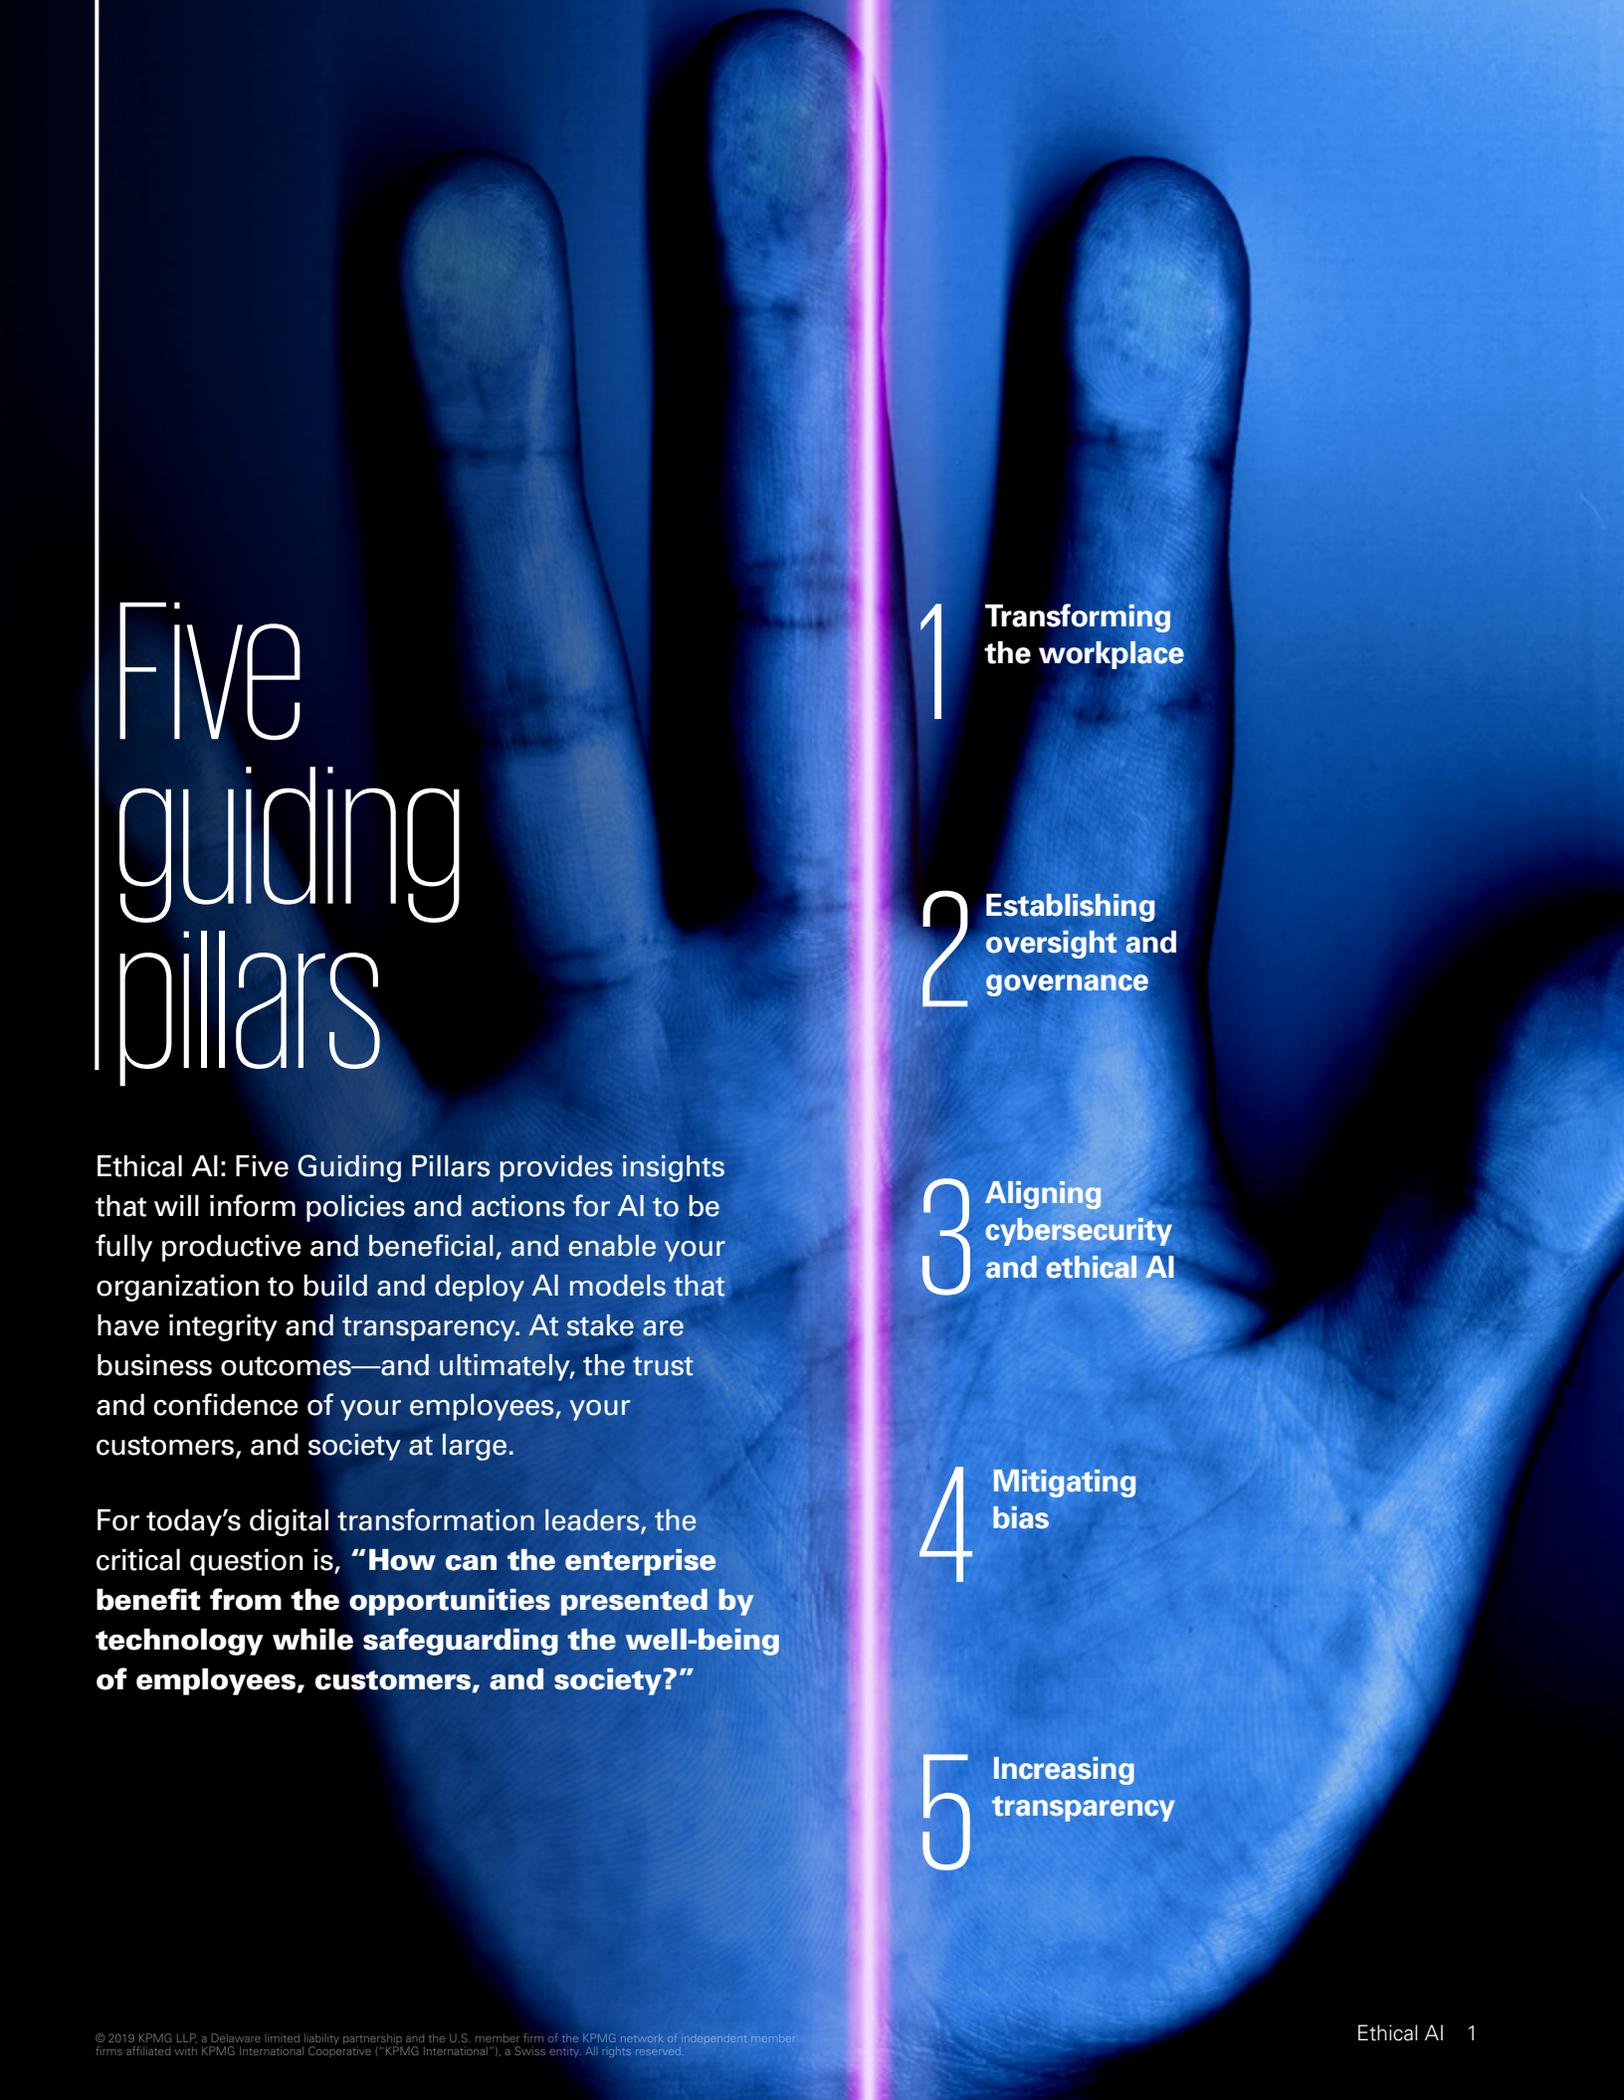
Todd Lohr oversees the KPMG Digital Lighthouse for the U.S. In this role, he works with clients to accelerate innovation, drive speed, and ensure global scale for data, analytics, and AI initiatives across the firm. He is a thought leader on Ethical Workforce Transformation around AI and has led numerous initiatives at Fortune 100 companies and government agencies to build frameworks and sustainable processes that help govern, secure, and control AI systems across the enterprise, helping to build and deploy AI ecosystems that have integrity and transparency. In addition to his practice responsibilities, Todd is a regular presenter at conferences and a guest lecturer at universities.



Traci Gusher

Principal, Innovation and Enterprise Solutions,
Artificial Intelligence

Traci Gusher leads the firm's vision and roadmap for its Artificial Intelligence solutions, including KPMG Ignite (KPMG's Global AI platform), and other data signals and intelligence platforms. She is an experienced executive having led teams who use AI, advanced analytics, and big data technologies to help clients drive growth, cut costs, and eliminate risks. Traci is passionate about "AI for Good," and works with her clients to help them see the overall good that machine learning can bring to customers, employees, and society, while helping them mitigate the risks associated with their algorithms. Traci is a frequent speaker on the topic of becoming an AI-driven enterprise.



Five guiding pillars

Ethical AI: Five Guiding Pillars provides insights that will inform policies and actions for AI to be fully productive and beneficial, and enable your organization to build and deploy AI models that have integrity and transparency. At stake are business outcomes—and ultimately, the trust and confidence of your employees, your customers, and society at large.

For today's digital transformation leaders, the critical question is, **“How can the enterprise benefit from the opportunities presented by technology while safeguarding the well-being of employees, customers, and society?”**

1 Transforming
the workplace

2 Establishing
oversight and
governance

3 Aligning
cybersecurity
and ethical AI

4 Mitigating
bias

5 Increasing
transparency



The impact of AI

Many children alive today will grow up to be AI natives. The presence of intelligent machines in their lives—at work, at play, as consumers and as citizens—will be as natural to them as digital experiences now are for us. They'll tap into the intelligence of AI as fluidly as we look for answers on the internet. And they won't be able to grasp that we once spent a great deal of our time on repetitive tasks that seem made for machines.

The level and speed of disruption and change that business leaders and society are facing in the advance of AI is without precedent in human history. Every other transformative technology emerged slowly, so that society had time to adjust and establish guardrails.

We are living in an era where we can put insights from AI to work on an extraordinary range of societal and scientific challenges. On a business level, AI propels new product development, enables epic customer experiences, and changes the nature of work itself. But to be truly successful, it must be deployed responsibly.

Corporate responsibility is not a new mission, but it has become a more complicated one as machine learning assumes a larger role in how work is done. As a matter of urgency and obligation, enterprises must consider how to address the immense societal impact that will come as work and decision-making change profoundly.

To create an ethical enterprise and sustain it into the future, AI must be governed and controlled in a meaningful way. This is what we mean by ethical AI.

Five guiding pillars

1 Transforming the workplace

The most immediate challenge facing leaders is the disruption of the workplace—the massive change in roles and tasks that define work, and the rise of powerful analytics and automated decision-making.

There is no doubt the job landscape will change, but according to our KPMG 2019 CEO Outlook, one area where U.S. CEOs are decidedly more optimistic this year is the impact of AI on jobs.

73% of those surveyed believe that AI and robotics will create more jobs than it eliminates, compared with 52 percent last year.¹

While a lot of attention has focused on the potential negative impact of automation technologies on jobs, we believe that this focus detracts from what is actually important: building the skills that companies will need to succeed and drive long-term value in a digitized world.

The workforce of the future calls for a new approach in business, one that is completely employee-centered and transparent.

KPMG's Chief Culture Officer, Claudia Saran, believes that properly integrating technology and the human workforce not only helps a company's financial results, but serves to strengthen a company's culture. For example, 92 percent of U.S. CEOs say they want their employees to feel empowered to innovate without worrying about negative consequences for them if their initiatives fail.¹ "One of the enablers of an innovative culture is a high comfort level with disruptive technologies," says Saran. "Those same technologies can also be utilized to help the organization learn from its innovation efforts – both successes and failures."

ACTION: Organizations can prepare for these changes now by helping employees adjust to the role of machines in their jobs. Consider partnering with academic institutions or other knowledge-leading organizations to create or pilot the programs and curricula that will directly address the need for new skills. On-the-job education through programs designed to prepare workers for change is a critical step that leading organizations have already begun to take. Such initiatives can help organizations hire efficiently, and it will help students find an in-demand career path.

At the same time, leaders need to acknowledge that workplace transformation is a change management challenge: Clearly communicate the changes and benefits of adoption, down to the task level, and cultivate new roles and tasks that add value for work and the business. Management must acknowledge and recognize individual needs and determine how best to help them pursue training and development that is aligned to their skills and interests. Consider creating surveys that check the pulse on workforce readiness and come up with a plan to close those capability gaps. Develop specific training and education initiatives and clearly communicate those efforts to target audiences will be key factors in success.

“

How do we assess the trade-off between the displacement effect of automation and the productivity improvement that's going to lead to greater employment and higher wages?

We found that when tasks are displaced through automation, there are also new tasks that get created. It becomes a race between the automation that's occurring and the new tasks being created. It's the balance between the two. ”

— Martin Fleming, Chief Economist, IBM, from the [KPMG Symposium on Artificial Intelligence](#), June 2019



The AI Compass: guiding success



“**AI-driven enterprises know where and when to use AI. They have an *AI compass* that helps point them in the right direction for governance, explainability, and value.**”

— Traci Gusher, Principal, Innovation and Enterprise Solutions, Artificial Intelligence

What are the qualities of AI-enabled organizations? First, they’re building an “AI compass” formed from a framework of algorithmic control that points them in the right direction in their business decisions. Driving growth, cutting costs, eliminating risks, and other business objectives are still priorities, but these AI leaders go further: They’re focused on the overall good that machine learning can bring to customers, employees, and society, and they devote resources and operational capabilities to the mitigation of risks associated with their algorithms.

Use of this “compass” is most evident in their approach to the workplace. The real AI enterprise knows that having strong data scientists and engineers alone will not set the needle in the right direction. The entire equation for transformation doesn’t happen until intentional efforts are made to transform the entire workforce.

2 Establishing oversight and governance

Emerging initiatives point to the potential for new regulation and/or oversight of AI. Consider two examples:

- The General Data Protection Regulation (GDPR), put forth and enacted by the European Union in 2018, is a major piece of legislation which, among other initiatives in the region, establishes guidelines for the use of AI. Declarations and guiding principles to safeguard the well-being of people as AI develops are already being formulated and may influence future regulations in the E.U.²
- The establishment in February 2019 of the American AI Initiative includes five areas of emphasis: investing in AI R&D; unleashing AI resources (data, models and computing capabilities, talent, and research); setting governance standards; building the AI workforce; and engagement to protect the U.S.'s advantage in AI.

These developments reinforce our view that businesses should launch education initiatives now, ahead of the deeper workforce disruption from AI to come. Ultimately, the lead in educating, training, and managing the AI-enabled workforce rests with business—and the sooner leaders set forth on this journey, the more influence they will have on evolving initiatives and regulations.

ACTION: Establish clear enterprise-wide policies now about the deployment of AI, including the use of data and standards of privacy. Taking this important step will keep business ahead of the curve—and it may help influence the direction of legislation, as policymakers look to successful organizations as role models. Leaders at the World Economic Forum (WEF) in Davos in August 2019 advocated a set of guidelines for responsible AI deployment that includes an “AI Toolkit” to help board members understand leading practices in AI deployment and methods for coping with bad PR and upset customers when AI problems occur.

Governing data and AI. Every business and organization—whether it's a global enterprise, local government, federal agency or NGO—needs to understand the data and frameworks that underlie their AI so they can use that knowledge to engender trust in their stakeholders. Brands need to revere data and establish strong best governance practices. These groups or their partners will also need to monitor outcomes continuously to be sure they are fair and accurate—and that they remain true to the original objectives.

Leaders in business and government can establish lines of ownership for the build and deployment of complex algorithms.

2. Oversight and governance (continued)

A recent report by KPMG introduces the AI in Control framework of methodologies and toolsets that can address the inherent risks and ethical issues in AI throughout its development cycle.³ Why is this so important?

The algorithms that are changing—and will continue to change our world—learn continuously based on the data they receive. The initial training data that underlies technologies like machine learning and deep learning have an obvious influence on the results. The data has to be accurate and relevant to the goal. But machine learning systems continue to absorb new or unseen data—and they continue to make decisions—based on changing environments. It's a continuous feedback loop of learning.

ACTION: Build algorithms based on a strong ethical framework. This is a core imperative for leaders charged with deploying AI responsibly. Leaders in business and government should establish lines of ownership for the build and deployment of complex algorithms. Leaders can take steps now to establish lines of accountability. Such steps may include hiring ethics officers into the C-suite, creating steering committees that include diverse leaders from both the technology and business sides, or bringing on a partner with an established framework for the governance and control of AI.

The algorithms that are changing—and will continue to change our world—learn continuously based on the data they receive.



3 Aligning cybersecurity and ethical AI

Autonomous algorithms prompt concern about cybersecurity risks as their influence on the business increases. Adversarial attacks that poison algorithms by tampering with training data can potentially compromise privacy and create bias—sully the user experience and damaging intellectual property, brand reputation, or other key business assets.

Integrating humans and machines in a way that makes the most of both skill sets is the big challenge enterprises now face. This is one reason why governing and controlling machine learning systems is an urgent priority. Companies recognize that mitigating algorithmic risk is a business and brand imperative. The arrival of ethics and responsibility officers in the C-suite at big brands like Salesforce and Uber is a signpost. But the need goes much deeper, as increasingly sophisticated machine learning systems influence people's behavior and expectations.

KPMG research conducted by industry insiders shows that potential threats to security are weighing heavily on industry leaders as they advance AI adoption.⁴

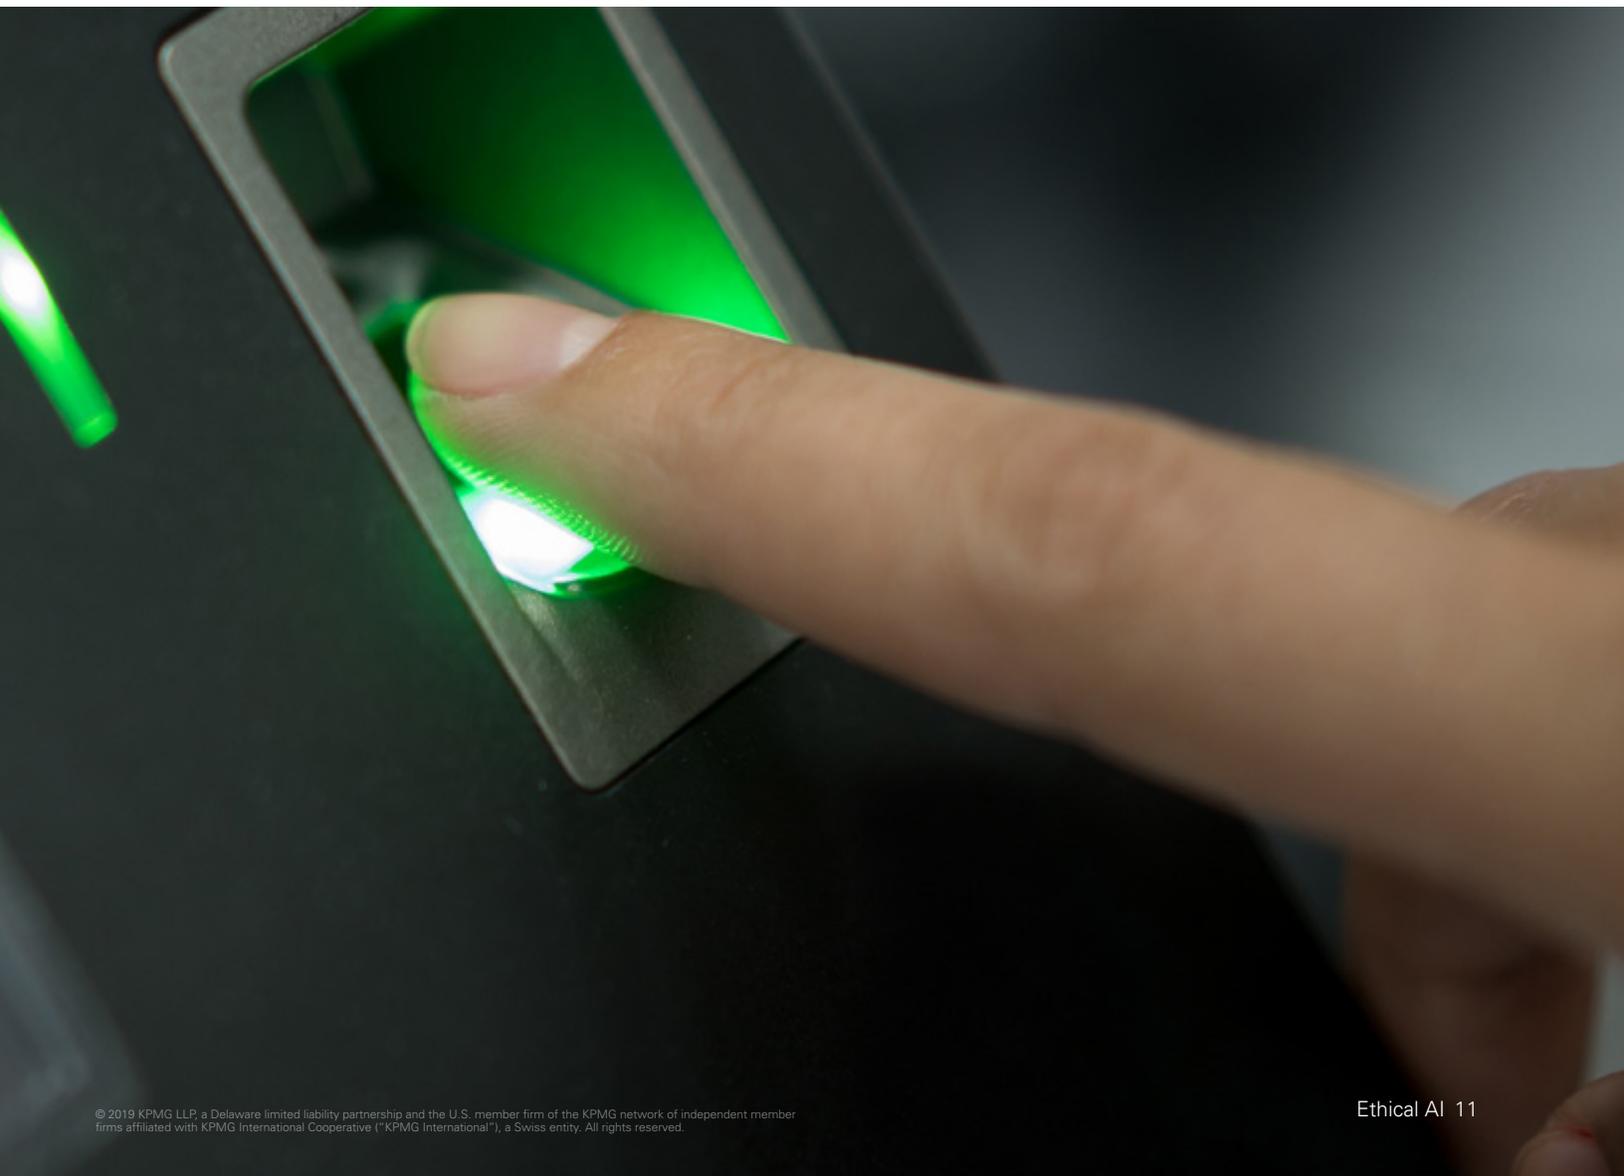
When thinking of their primary ethical concerns when it comes to AI tech, Healthcare (75%) and Finance leaders (72%) are especially concerned about threats to security, as are most Retail (61%), Transportation (62%) and Technology (59%) insiders.

72% of U.S. CEOs agree this year that strong cybersecurity is critical to engender trust with their key stakeholders, compared with 15% last year.¹

ACTION: Build strong security into the creation of algorithms and the governance of data.

When an “ethical compass” is in the DNA of algorithmic development, the overall integrity of the model will be stronger. These aspects contribute to building security in AI:

- Understand the context and intended purpose under which the model was developed
- Identify who trained the algorithms
- Know the provenance of the data and any changes made to it
- Understand how the models were (and are) served and protected
- Maintain continuous review and confirmation of an algorithm’s effectiveness and accuracy.



4 Mitigating bias

Leaders and regulators want to understand the workings of sophisticated, autonomous algorithms so they can take steps to eliminate unfair bias over time as they continue to evolve. That's why understanding which attributes in the training data influence a model's predictions has become very important. When training lapses for long periods, a model can lose integrity: bias can creep in, and issues such as adversarial attacks can compromise what these models learn.

Every leader should have a moral imperative to mitigate bias by governing AI along its entire lifecycle—from its ideation and build to its continuing evolution—and then take new steps to manage and guide an increasingly diverse workforce as the nature of work changes. Many leaders are taking specific actions now before new regulations demand compliance with new standards.

Controlling sophisticated algorithms is at the heart of ethical AI. Algorithms won't be trusted if they're not fair – and to be fair, they need to be designed and built to be as free from bias as possible, and they need to maintain their fairness as they evolve. Attributes used to train algorithms must be relevant, appropriate for the goal, and allowed for use. In some instances, however, personal information is

relevant to the model. That is why many firms have hired an AI ethicist, and it's why transparency is so crucial for AI technologies to evolve and brands to thrive.

According to Harvard Business Review,⁴ we need to shift from a culture of automation to augmentation. Authors Satell and Sutton claim that artificial intelligence works best not as some sort of magic box used to replace humans and cut costs, but as a force multiplier that used to create new value. By making AI more explainable, auditable, and transparent, we can make our systems more fair, more effective and more useful.

ACTION: Ensure the goal and purpose of critical algorithms are clearly defined and documented.

Verify that the design aligns with principles, standards, and guidelines, corporate values and ethics, compliance, and security and quality standards. The best way to identify and avoid bias may lie with teams charged with leading the training of algorithms. It's also important to set up an independent review of critical models where bias could produce an adverse social impact. Without actions like these the very opportunities and benefits of AI may be tarnished or even ruined.

By making AI more explainable, auditable, and transparent, we can make our systems more fair, more effective and more useful.

5 Increasing transparency

Some leaders are proposing universal standards for fairness as AI and machine learning advance. The goal would be transparency—and, above all, clarity—for citizens and consumers who are confused about the data and information they provide in digital (and analog) experiences. If business can firmly establish its own policies and practices now—and demonstrate success with them—it will inform standards at a national or federal level.

The power and promise of AI can only be fully unleashed by our understanding and control of its build and actions—actions that, depending on its integrity, can either propel us forward or set us back. This is why companies need to establish an overall management policy for AI, with a focus on unleashing the power of AI technologies responsibly.

Businesses—especially some of the global technology companies—are creating clear messaging and user interfaces so people can easily understand and decide how to let brands deliver service. This should be the “default setting” for any brand’s focus on building and bolstering trust.

ACTION: Create “contracts of trust.” Give customers the clarity and information that they they want and need. That simple, basic guiding principle is what brands must do. The tricky part is being simple and clear about complex issues, explaining how data is being used and why it influences a customer’s experience. Let the public know how you’re being transparent and what decisions about personal data will mean. You’re giving the customer what they want—the power to opt-in or opt out—and you’re showing them the benefit they derive by providing their permission. Think of your handling of personally identifiable information and data as a “contract of trust.”



The management policy should not be limited to the build and control of algorithms—it should extend to how the brand communicates with its customers and the public at large.

Conclusion: Ethical AI imperatives

Fear of lost revenue, reputational damage or regulatory fines should *not* be the driving force behind your decision to embark on an ethical transformation with your AI, according to Forrester Research.⁵

Love for your company, mission, customers, and colleagues is what gives your company authentic values. According to Forrester, the ethical transformation with your AI is an opportunity to embrace change, where your stakeholders, rather than your board, drive the company's ethics.

We've found that many business leaders are at a loss to translate talk into concrete governance and control in a meaningful way.

Almost no organization is truly prepared at this moment to grapple with—much less solve—the profound questions posed by AI technologies. But they can take specific steps now.

Taking action

Ethical AI is about taking action. KPMG has distilled the actions necessary to point an organization toward a “true north” of corporate and civil ethics around AI.



Prepare employees now: Partner with academia or other knowledge-leading organizations to help create or pilot the programs that will directly address the need for new skills and help employees adjust to the role of machines in their jobs.



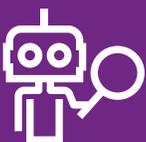
Develop strong oversight and governance: Establish clear policies about the deployment of AI, including the use of data, standards of privacy, and governance of leading practices.



Align cybersecurity and ethical AI: Build strong security into the creation of algorithms and the governance of data.



Mitigate bias: Ensure the goal and purpose of critical algorithms are clearly defined and documented. Attributes used to train algorithms must be relevant, appropriate for the goal, and allowed for use.



Increase transparency: Create “contracts of trust.” Let the public know how you are being transparent and what decisions about their personal data will mean to them.

KPMG Trusted Enterprise

Establishing the mechanisms to govern, secure, and control AI systems is a mission-critical part of developing and deploying AI ethically. KPMG can help your business take action across these five critical pillars of Ethical AI so you can fully benefit from the opportunities presented by these approaches, while minimizing risk.

Those organizations that bring transparency, data protection, control, and privacy rigor to the handling of data and AI are rewarded with increased trust – and the brand loyalty and increased market share that goes with it.

Contact KPMG to learn more about how we can help you benefit from the opportunities presented by technology—while safeguarding the well-being of employees, customers, and society—and become a Trusted Enterprise.

References

- 1 [KPMG 2019 U.S. CEO Outlook](#)
- 2 Declaration on Ethics and Data Protection in Artificial Intelligence, 40th International Conference of Data Protection and Privacy Commission, October 2019
- 3 [Controlling AI: The imperative for transparency and explainability](#), June 2019
- 4 KPMG research conducted with 750 industry insiders in October 2019
- 5 Harvard Business Review, "We Need AI That Is Explainable, Auditable, and Transparent," Greg Satell, Josh Sutton, October 28, 2019
- 6 Forrester Research, "Make Your Digital Transformation An Ethical One," July 8, 2019

Contact

Todd Lohr

Principal, Advisory,
KPMG Digital Lighthouse
Network Leader
tlohr@kpmg.com
+1 312 665 1684

Traci Gusher

Principal, Advisory,
Innovation and
Enterprise Solutions
tgusher@kpmg.com
+1 267 256 1617

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates and related entities.

The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act upon such information without appropriate professional advice after a thorough examination of the particular situation.

kpmg.com/socialmedia



© 2019 KPMG LLP, a Delaware limited liability partnership and the U.S. member firm of the KPMG network of independent member firms affiliated with KPMG International Cooperative ("KPMG International"), a Swiss entity. All rights reserved. Printed in the U.S.A. The KPMG name and logo are registered trademarks or trademarks of KPMG International.