



# Putting principles into practice: How we approach responsible AI at Microsoft



# Table of Contents

- Embracing the transformative potential of AI..... 3**
- Microsoft AI guiding principles..... 4**
  - Fairness.....4
  - Reliability and safety .....5
  - Privacy and security .....5
  - Inclusiveness.....6
  - Transparency.....6
  - Accountability .....7
- Recommendations for how to implement responsible AI ..... 8**
  - Establishing an AI governance system.....8
    - Understanding Microsoft’s governance model ..... 9
    - Governance in action at Microsoft: sensitive use case framework ..... 11
  - Establishing design principles and guidelines in engineering ..... 13
  - Engineering tools for responsible AI ..... 14
    - Fairness .....14
    - Reliability and safety ..... 15
    - Security and privacy ..... 15
    - Inclusiveness..... 16
    - Transparency ..... 16
    - Accountability ..... 16
  - Engaging with external stakeholders ..... 17
    - Participating in industry coalitions ..... 17
    - Providing a perspective on policy..... 17
    - Addressing future labor and workplace needs..... 18
    - Contributing solutions to societal challenges..... 19
      - AI for Accessibility..... 19
      - AI for Earth..... 19
      - AI for Humanitarian Action ..... 20
      - AI for Cultural Heritage..... 20
      - AI for Health..... 20
- Working together for a better future .....21**
- Learn more today.....21**





## Embracing the transformative potential of AI

AI is the defining technology of our time. It is already enabling faster and more profound progress in nearly every field of human endeavor and helping to address some of society's most daunting challenges—like providing remote students with access to education and helping farmers produce enough food for our growing global population.

At Microsoft, we believe that the computational intelligence of AI should be used to amplify the innate creativity and ingenuity of humans. As both a supplier of AI technology and a company that increasingly uses AI to run its business, our vision for AI is to empower every developer to innovate, empower organizations to transform industries, and empower people to transform society. To that end, we are not only focused on helping organizations take full advantage of AI, but we are investing heavily in AI for Good programs that provide technology, resources, and expertise to empower those working to create a more sustainable, safe, and accessible world.

But as with all great technological innovations in the past, the use of AI technology will have broad impacts on society, raising complex and challenging questions about the future we want to see. AI will have implications on decision-making across industries, data security and privacy, and the skills people need to succeed in the workplace. As we look to this future, we must ask ourselves: How do we design, build, and use AI systems that create a positive impact on individuals and society? How can we best prepare workers for the impact of AI? How can we attain the benefits of AI while respecting privacy?

Stakeholders across academia, government, and the private sector are considering these questions as they develop and implement AI. Laws and regulations are emerging, but at the same time organizations are finding the need to create internal policies and practices to guide their specific AI efforts.<sup>i</sup> While organizations can't predict the future just yet, it's our responsibility to make a concerted effort to anticipate and mitigate the unintended consequences of the technology we release into the world through deliberate planning and continual oversight.

In this document, we want to share what we're learning in our own journey in hope of providing a useful perspective for other organizations navigating similar challenges. We started by establishing core principles to guide our development and use of AI products and solutions. Using those principles as a foundation, we then developed an internal governance system to help us engage with AI responsibly and shared resources, guidance, and tools with our engineering teams. Finally, we are making concrete investments to solve key societal challenges by engaging with external stakeholders and investing in our "AI for Good" initiatives. We ultimately aim to foster responsible AI that lives up to our belief that AI can serve humanity and solve some of the most profound challenges facing our world.



## Microsoft AI guiding principles

At Microsoft, we've recognized six principles that we believe should guide AI development and use: fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability.

For us, these principles are the cornerstone of a responsible and trustworthy approach to AI, especially as intelligent technology becomes more prevalent in the products and services we use every day.

### Fairness

AI systems should treat everyone fairly and avoid affecting similar situated groups of people in different ways. For example, when AI systems provide guidance on medical treatment, loan applications or employment, they should make the same recommendations to everyone with similar symptoms, financial circumstances or professional qualifications.

We believe that mitigating bias starts with people understanding the implications and limitations of AI predictions and recommendations. While an AI system can provide useful suggestions, people should always be accountable for consequential decisions that affect others. It will be important to train people to understand the meaning and implications of AI decisions and supplement those decisions with sound human judgment.

When designing and building AI systems, developers should understand how bias can be introduced and how it can affect AI-based recommendations.

To help mitigate bias, they should use training datasets that reflect the diversity of society and design models in ways that allow them to learn and adapt over time without developing biases.

For example, consider an AI system designed to help employers screen job applicants. When trained on data from public employment records, this system might “learn” that most software developers are male and favor men over women when recommending candidates. We encountered a similar situation when partnering with a large financial lending institution to develop a risk scoring system for loan approvals. We trained an existing industry model using the customer’s data. When we conducted an audit of the system, we discovered that while it only approved low-risk loans, all approved loans were for male borrowers. The training data reflected the fact that loan officers historically favor male borrowers—and inspecting the system allowed us to identify and address that bias before the system was deployed.

To help them develop AI systems that treat everyone fairly, developers can leverage tools, methodologies, techniques, and other resources that help detect and mitigate biases.

#### Some considerations to promote fairness

1. Understand the scope, intent, and potential uses of the AI system
2. Attract a diverse pool of talent
3. Put processes and tools in place to identify bias in datasets and machine learning algorithms
4. Leverage human review and domain expertise
5. Research and employ best practices, techniques, and tools
6. Train AI systems with data that represents the world

See more on the [AI Business School](#)



## Reliability and safety

To build trust, it's critical that AI systems operate reliably, safely, and consistently under normal circumstances and in unexpected conditions. These systems should be able to operate as they were originally designed, respond safely to unanticipated conditions, and resist harmful manipulation. It's also important to be able to verify that these systems are behaving as intended under actual operating conditions. How they behave and the variety of conditions they can handle reliably and safely largely reflects the range of situations and circumstances that developers anticipate during design and testing.

We believe that rigorous testing is essential during system development and deployment to ensure AI systems can respond safely in unanticipated situations and edge cases, don't have unexpected performance failures, and don't evolve in ways that are inconsistent with original expectations.

After testing and deployment, it's equally important that organizations properly operate, maintain, and protect their AI systems over the lifespan of their use. Otherwise, AI systems can become unreliable or inaccurate if they are not maintained properly.

Ultimately, because AI should augment and amplify human capabilities, people need to play a critical role in making decisions about how and when an AI system is deployed, and whether it's appropriate to continue to use it over time. Human judgment will be key to identifying potential blind spots and biases in AI systems.

### Some considerations to promote reliability and safety

1. Understand your AI maturity
2. Develop processes for auditing AI systems
3. Provide detailed explanations of system operation
4. Design for unintended circumstances
5. Involve domain experts in design and implementation processes
6. Conduct rigorous testing during AI system development and deployment
7. Evaluate when and how an AI system should seek human input
8. Develop a robust feedback mechanism for user to report performance issues

See more on the [AI Business School](#)



## Privacy and security

As AI becomes more prevalent, protecting privacy and securing important personal and business information is becoming more critical and complex. With AI, privacy and data security issues require especially close attention because access to data is essential for AI systems to make accurate and informed predictions and decisions about people. AI systems must comply with privacy laws that require transparency about the collection, use, and storage of data and mandate that consumers have appropriate controls to choose how their data is used.

As new intelligent technology emerges, organizations must be prepared for emerging and evolving privacy and security threats. We were reminded of this responsibility in 2016 when we released a chatbot on Twitter called Tay. We taught Tay to learn unsupervised from interactions with Twitter users, so she could better replicate human communication and personality traits. However, within 24 hours users realized that she could learn and began to feed her bigoted rhetoric, turning her from a polite bot into a vehicle for hate speech.

### Some considerations to promote privacy and security

1. Comply with relevant data protection, privacy, and transparency laws
2. Protect AI systems from bad actors
3. Design AI systems with appropriate customer controls
4. Ensure your AI system maintains the anonymity and integrity of personal data
5. Conduct security and privacy reviews
6. Research and implement industry best practices

See more on the [AI Business School](#)

This experience taught us that while technology may not be unethical on its own, people do not always have good intentions and we must consider the human element when designing AI systems. We learned to prepare for new types of attacks that influence learning datasets, especially for AI systems that have automatic learning capabilities. To help ensure a similar experience does not happen again, we developed technology such as advanced content filters and introduced supervisors for AI systems with automatic learning capabilities.

## Inclusiveness

At Microsoft, we firmly believe everyone should benefit from intelligent technology, meaning it must incorporate and address a broad range of human needs and experiences. For the 1 billion people with disabilities around the world, AI technologies can be a game-changer. AI can improve access to education, government services, employment, information, and a wide range of other opportunities. Intelligent solutions such as real-time speech-to-text transcription, visual recognition services, and predictive text functionality are already empowering those with hearing, visual, and other impairments.

[Inclusive design practices](#) can help system developers understand and address potential barriers in a product environment that could unintentionally exclude people. By addressing these barriers, we create opportunities to innovate and design better experiences that benefit everyone.

## Transparency

Underlying the preceding values are two foundational principles that are essential for ensuring the effectiveness of the rest: transparency and accountability. When AI systems are used to help inform decisions that have tremendous impacts on people's lives, it is critical that people understand how those decisions were made. For example, a bank might use an AI system to decide whether a person is creditworthy, or a company might use an AI system to determine the most qualified candidates to hire.

A crucial part of transparency is what we refer to as intelligibility, or the useful explanation of the behavior of AI systems and their components. Improving intelligibility requires that stakeholders comprehend how and why they function so that they can identify potential performance issues, safety and privacy concerns, biases, exclusionary practices, or unintended outcomes. We also believe that those who use AI systems should be honest and forthcoming about when, why, and how they choose to deploy them.

### AI & GDPR

For more information on deploying AI while complying with Europe's GDPR, we encourage you to read our AI Insights feature [Artificial Intelligence and the GDPR Challenge](#)

### Some considerations to promote inclusiveness

1. Comply with laws regarding accessibility and inclusiveness
2. Use the [Inclusive Design toolkit](#)
3. Have people with disabilities test your systems
4. Consider leveraging commonly used accessibility standards

See more on the [AI Business School](#)

### Some considerations to promote transparency

1. Share key characteristics of datasets
2. Improve model intelligibility by using simpler models and explaining model behavior
3. Train employees on how to interpret AI outputs

See more on the [AI Business School](#)





## Accountability

The people who design and deploy AI systems must be accountable for how their systems operate. Organizations should draw upon industry standards to develop accountability norms. These norms can ensure that AI systems are not the final authority on any decision that impacts people's lives and that humans maintain meaningful control over otherwise highly autonomous AI systems.

Organizations should also consider establishing a dedicated internal review body. This body can provide oversight and guidance to the highest levels of the company on which practices should be adopted to help address the concerns discussed above and on particularly important questions regarding the development and deployment of AI systems. They can also help with tasks like defining best practices for documenting and testing AI systems during development or providing guidance when an AI system will be used in sensitive cases (like those that may deny people consequential services like healthcare or employment, create risk of physical or emotional harm, or infringe on human rights).

The need for accountability is particularly crucial with sensitive technologies like facial recognition. Recently, there has been a growing demand for facial recognition technology, especially from law enforcement organizations that see the potential of the technology for use cases like finding missing children. However, we recognize that these technologies could potentially be used by a government to put fundamental freedoms at risk by, for example, enabling continuous surveillance of specific individuals. We believe society has a responsibility to set appropriate boundaries for the use of these technologies, which includes ensuring governmental use of facial recognition technology remains subject to the rule of law.

While we believe that new laws and regulations are indispensable, we also recognize that they are not a substitute for the responsibility that needs to be exercised by businesses, governments, NGOs, and academic researchers engaging with AI. This is why, in July 2018, we announced that we would assess and develop principles to govern our work with facial recognition technologies. We anticipate these principles will evolve over time as we continue to learn and partner with customers, other tech companies, academics, civil society, and others on this issue.

### The Future Computed

To learn more about our six principles as well as the impact of AI on our future, please read our book, [The Future Computed](#).

Facial recognition technology highlights the importance of preparing for and remaining vigilant of shortcomings and unintended consequences with all emerging AI. We consider it a shared responsibility across the public and private sector to engage with AI responsibly. It's essential that we continue to foster open dialogue among businesses, governments, NGOs, academic researchers, and all other interested individuals and organizations.

### Some considerations to promote accountability

1. Set up internal review boards
2. Ensure your employees are properly trained to use and maintain the solution
3. Keep humans with requisite expertise in the loop
4. Put a clear system of accountability and governance in place

See more on the [AI Business School](#)



## Recommendations for how to implement responsible AI

Each organization will have its own guiding principles, but ultimately these principles need to be part of a larger responsible AI strategy to be effective. A holistic AI strategy should help you bring your principles to life both within your organization and beyond.

Within your organization, we recommend establishing a system for internal oversight, or “governance,” to provide guardrails for how AI solutions are designed, trained, deployed, and monitored. This system could be tailored to your organization’s unique characteristics, culture, guiding principles, and level of engagement with AI. Beyond internal governance, we encourage you to consider how to foster responsible and trustworthy AI in your industry and throughout society. We all must work together to maximize AI’s potential for positive change.

See the Microsoft AI Business School for [more recommendations](#) on how to leverage governance and external engagements in your organization.

### Establishing an AI governance system

AI is more than just a new technology. Because of the far-reaching impacts it can have in an organization, businesses and governments alike are seeing a growing need to create internal governance systems to help them engage with AI responsibly.<sup>1</sup>

As discussed above, the first step for any AI governance system is to establish guiding principles for the organization. This is especially important because laws and regulations around ethics and technology are relatively nascent.

With guiding principles providing a foundation, governance systems are generally tasked with objectives such as:

- Developing and enforcing policies, standards, and best practices for testing, documenting, and managing AI systems
  - Policies may include rigorous testing and pilots before deployment, conducting regular audits of AI systems, monitoring AI systems for model drift and decay over time, processes for holding employees accountable, or standards for AI security and privacy.
- Providing advice on ethical concerns
- Training employees who design or use AI systems to understand models and data sources, identify bias in datasets or AI recommendations, and report performance issues.
- Creating guidelines on how to safely operate and monitor an AI system, and create processes to ensure they are being followed throughout the AI system’s entire lifespan



For any type of governance system to be successful, it must report to top decision-makers and have the financial and human resources, authority, and processes to affect real change throughout the organization. No matter what model you choose, it will be important to be prepared to evolve over time.

## Understanding Microsoft's governance model

In keeping with our culture of integrity and trust at Microsoft, we knew it was essential to create a governance structure that would enable responsible AI innovation from the ground up. From the senior leadership team to developers, field sellers and beyond, everyone is empowered to play a role in fostering responsible AI capabilities.

While we knew that engaging with AI would require careful consideration and oversight, we didn't have a perfect responsible AI governance system on day one. In fact, our governance system continues to evolve to this day. And that's as it should be. A governance system should be agile to the changing nature of technology and the business. From the beginning, we created a tight loop between engineering and our responsible AI advisory committee called Aether (AI, Ethics, and Effects in Engineering and Research). Engineering leadership serves on the committee and engineering practitioners are included in Aether working groups. We recommend forming an advisory committee as soon as you invest in AI engineering (even if it's a year or more before products hit the market).

Initially, to govern responsible AI across the company, the Aether committee reviewed and advised on sensitive AI use cases as they arose. We have built upon this early case review model to add company-wide rules that set out practices designed to uphold our principles, education and training materials, a Responsible AI Champs program, and even tools that plug directly into sales and engineering workflows. Our governance structure today uses a hub-and-spoke model to provide the accountability and authority to drive initiatives while also enabling responsible AI policies to be implemented at scale.

### Centralized governance

The **Senior Leadership Team** is ultimately accountable for the company's direction on responsible AI, setting the company's AI principles, values, and human rights commitments. Building off our culture of integrity and trust, this group is the final decision-maker on the most sensitive, novel, and significant AI development and deployment matters.

Our commitment to responsible AI governance is administered, implemented, and maintained by the **Office of Responsible AI**. The office works with stakeholders across the company to develop and maintain our governance framework, define roles and responsibilities for governing bodies, implement a company-wide reporting and decision-making process, and orchestrate responsible AI training for all employees.

The Office of Responsible AI has four key functions:

- Internal policy: Setting the company-wide rules for enacting responsible AI, as well as defining roles and responsibilities for teams involved in this effort.
- Enablement: Readiness to adopt responsible AI practices, both within our company, and among our customers and partners.
- Case management: Review of sensitive use cases to help ensure that our AI principles are upheld in our development and deployment work.
- Public policy: Help to shape new laws, norms, and standards that will be needed to ensure that the promise of AI technology is realized for the benefit of society at large.

The **Aether Committee** serves an advisory role to the senior leadership and the Office of Responsible AI on questions, challenges, and opportunities with the development and fielding of AI technologies. Aether also

provides guidance to teams across the company to ensure that AI products and services align with our AI principles. The committee brings together top talent in technology, ethics, law, and policy from across Microsoft to formulate recommendations on policies, processes, and best practices.

The Aether Committee has six **working groups** that focus on specific topics, grounded in our AI principles. The working groups play a key role in developing tools, best practices, and tailored implementation guidance related to their respective areas of expertise. Learnings from the working groups and main committee have resulted in Microsoft developing new policies, and in some cases either declining or placing limits on specific customer engagements where AI-related risks were high.

---

#### AETHER working groups:

- Bias and Fairness of AI systems
  - Engineering Practices for AI
  - Human-AI Interaction
  - Intelligibility and Explanation of AI recommendations
  - Safety and Robustness
  - Sensitive Uses of AI
- 

Together, Aether and the Office of Responsible AI work closely with our engineering and sales teams to help them uphold our AI principles in their day-to-day work. An important hallmark of our approach to responsible AI is having this ecosystem to operationalize responsible AI across the company, rather than a single organization or individual leading this work. Our approach to responsible AI also leverages our process of building privacy and security into all products and services from the start.

### Decentralized functions

Enacting responsible AI at scale across an organization relies on a strong network across the company to help implement organization-wide rules, drive awareness, and request support on issues that raise questions about application of our AI principles.

Our network includes **Responsible AI Champs**, employees who are nominated by their leadership teams, from within key engineering and field teams to serve as responsible AI advisors (in addition to their full-time roles). Rather than devising the role as a policing function, they serve an advisory role to help inform decision-makers. The Responsible AI Champs have five key functions:

- Raising awareness of responsible AI principles and practices within teams and workgroups
- Helping teams and workgroups implement prescribed practices throughout the AI feature, product or service lifecycle
- Advising leaders on the benefit of responsible AI development – and the potential impact of unintended harms
- Identifying and escalating questions and sensitive uses of AI through available channels
- Fostering a culture of customer-centricity and global perspective, by growing a community of Responsible AI evangelists in their organizations and beyond

To develop and deploy AI with minimal friction to engineering practices and customers, we are investing in patterns, practices, and tools. Some **engineering groups** have assembled teams to help them follow the company-wide rules and accelerate the development of implementation patterns, practices, and tools.

The final and perhaps most important part of our approach to responsible AI is the role that **every employee** plays with support from their managers and business leaders. Responsible AI is a key part of mandatory employee training and we have released additional educational assets that enable employees to delve deeper into areas of responsible AI. We also have numerous responsible AI development tools to enable our employees to develop responsibly. With these resources, all our employees are empowered to advance the company's important work with AI and, at the same time, they are responsible for upholding our responsible AI principles and following the company-wide practices we have adopted in pursuit of that end.

## Policies and procedures

To help every Microsoft employee live up to our commitment to developing and deploying responsible AI, we have created principles, a sensitive use framework, and company-wide rules to help employees develop a better understanding of the company's commitment with respect to AI development and deployment.

We expect every Microsoft employee to:

- Develop a general understanding of our AI principles
- Report and escalate sensitive uses
- Contact their Responsible AI Champ when they need guidance on responsible AI.

## Responsible AI Standard

In order to implement responsible AI practices, the policy requirements, procedures, and tools need to be tightly embedded with the AI development lifecycle an organization already uses. In a recent study conducted by Microsoft, researchers found that there are consistent stages in the AI development process, however, there is not a consistent or common set of tools that individuals use across the discipline<sup>1</sup>. In order to assist teams in implementing policies and procedures, and to further operationalize our approach to responsible AI, Microsoft developed the Responsible AI Standard, which outlines the steps teams must follow to support the responsible development and deployment of AI systems. A key part of the Standard is organized by the Responsible AI Lifecycle (RAIL), helping to guide engineering teams through responsible AI considerations organized by stage of the AI development lifecycle. We're still in learning mode with the Responsible AI Standard as we pilot it in different engineering and sales teams across the company, and we look forward to sharing our learnings with you in time.

## Governance in action at Microsoft: sensitive use case framework

Microsoft leverages a "sensitive uses" review trigger to help our internal and customer-facing teams identify when specific use cases need additional guidance. This is one of the early steps in our responsible AI governance process. We've found that it's highly advantageous to start identifying, assessing, and mitigating responsible AI considerations early in solution design, rather than only as a compliance check at the end of a project.

Per our responsible AI governance documentation, an AI development or deployment scenario is considered a "sensitive use" if it falls into one or more of the following categories:

1. **Denial of consequential services:** The scenario involves the use of AI in a way that may directly result in the denial of consequential services or support to an individual (for example, financial, housing, insurance, education, employment, healthcare services, etc.).
2. **Risk of harm:** The scenario involves the use of AI in a way that may create a significant risk of physical, emotional, or psychological harm to an individual (for example, life or death decisions in military, safety-critical manufacturing environments, healthcare contexts, almost any scenario involving children or other vulnerable people, etc.).
3. **Infringement on human rights:** The scenario involves the use of AI in a way that may result in a significant restriction of personal



<sup>1</sup> Microsoft, "Software Engineering for Machine Learning: A Case Study," 2019.



freedom, opinion or expression, assembly or association, privacy, etc. (for example, in law enforcement or policing).

We train our employees to use this framework to determine whether an AI use case should be flagged for further review—whether they’re a seller working with a customer or someone working on an internal AI solution. We also train our Responsible AI Champs for their role as liaison between employees and central governance teams.

## Microsoft sensitive use case review process

### Identification

If an employee identifies that a use case falls into one of these three categories, they report it via a central submission tool and it’s routed to their local Responsible AI Champ—an individual who is responsible for driving awareness and understanding of the company’s responsible AI policies, standards, and guidance.

### Assessment

The Responsible AI Champ, working with the Office of Responsible AI and the Microsoft team involved in the use case, investigates the case to gather the relevant facts, follows a guided process to assess the impact of the proposed system on individuals and society, and reviews past cases to determine if guidance already exists for a similar scenario. If previous guidance does not exist, or if the case requires additional expertise and evaluation, the Responsible AI Champ presents the case to the Sensitive Uses Working Group of the Aether Committee (AI, Ethics, and Effects in Engineering and Research).

### Mitigation

The Sensitive Uses Working Group deliberates with a diverse group of experts to provide insight and recommendations for how to address the risks associated with the particular use case. If further escalation is needed, cases can rise all the way up to the Aether Committee itself, which directly advises Microsoft senior leadership. Ultimately, decisions on novel, high-impact cases are made by company leadership.

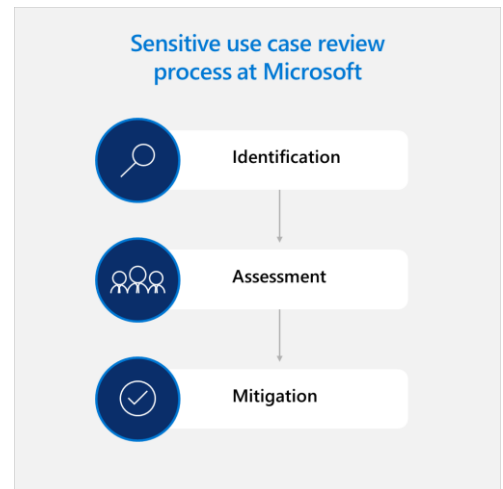
When reviewing sensitive use cases, we recommend bringing together a diverse group of people with varied backgrounds and expertise. It’s also important to create an inclusive space where everyone feels comfortable sharing their ideas and perspectives.

Once the case has been reviewed, the Responsible AI Champ works with the Office of Responsible AI to provide advice to the project team on mitigation strategies that align to our responsible AI practices and principles. These mitigation strategies could include technical approaches, employee training and governance approaches, or alterations to the scope of the project. At times, our teams have been advised to not proceed with certain projects because we were unable to deliver them in a way that upholds our principles.

## Responsible AI governance in action: case study

To better illustrate this process, let’s review a real-world case study where a customer came to us with a potentially sensitive use of AI.

We were approached by a law enforcement agency to develop a facial recognition system to augment existing identity verification methods. The scenarios included using facial recognition to check drivers’ identities during



traffic stops, to expedite the check-in process in prisons, and to verify prisoners' identities while moving through the facility. These three use-cases were submitted through the central intake tool for responsible AI review.

Through many thoughtful discussions, the Responsible AI Champ worked closely with the Office of Responsible AI, the account teams, and the customers to assess the risks. It was determined that all three cases needed to be escalated to the Aether Sensitive Uses Working Group for further input, as they touched on one or more of the sensitive uses of AI.

After careful consideration, we determined that we would not support the patrolling scenario to identify "persons of interest," during traffic stops. As the state of the technology and the broader ecosystem were not sufficiently mature enough to mitigate the harmful consequences for when the technology performs imperfectly, this scenario was considered a premature use case. In addition, attempting to identify individuals in uncontrolled environments can infringe on human rights, resulting in improper arrests due to misidentification. Studies have shown that AI is more likely to mistake the identities of women and minorities, which could also lead to those populations being disproportionately detained.<sup>1</sup> We explained the issues to the customer, and they decided not to pursue that scenario.

For the in-facility use cases, we determined we could support the design and development of a proof of concept (POC), with safeguards in place to ensure appropriate human control over the solution, and a bi-directional feedback loop between the customer and Microsoft could be established. It was also important that the customer implement a training program for personnel interacting with the solutions, and that the customer would reengage with Microsoft on deployments beyond these supported scenarios.

## The evolution of responsible AI governance

AI is still a relatively new field, so it should come as no surprise that the processes around it are evolving rapidly as well. Going forward, we plan on refining our governance policies as we invest further in AI, and we recommend other businesses do the same. Every organization will need to customize its review process based on its own unique AI ambitions, needs, and maturity, but hopefully our process can serve as a helpful starting point.

Next, let's review some guidelines and tools available to developers to help them build and deploy AI responsibly.

## Establishing design principles and guidelines in engineering

If you're developing, implementing, or managing AI internally, we recommend that you translate your guiding principles into actionable guidance for technical employees.

Microsoft is on this journey as well. For years, we have been working with other companies and institutions to help developers everywhere build and deploy AI responsibly. We also leverage open-source tools and look to leading organizations like Partnership on AI (PAI) for best practices, industry standards, and guidelines. By leveraging practical guidance, hopefully you and your team don't have to develop your approach from scratch.

First, we recommend that you provide developers with detailed and thorough **standard guidance**. Consider developing checklists that map to the phases of an AI project, from data acquisition to model development and testing.

Then, because AI systems vary widely in terms of purpose and operation, you might find it necessary to create or reference **guidelines for specific AI technologies**.

- For example, in May of 2019 we published a paper called [Guidelines for Human-AI Interaction](#), which includes 18 generally applicable guidelines to help developers design responsible and human-centered AI

systems. These guidelines synthesize 150+ design recommendations from over 20 years of research in academia and industry. They suggest best practices for how AI systems should behave upon initial interaction, during regular interaction, when they're inevitably wrong, and over time.

- We've also released guidelines to help developers build [responsible conversational AI](#). Like any technology, virtual assistants and bots can pose a significant risk when developed and deployed improperly. This is especially true when they are helping people navigate information related to employment, finances, or health. We created these guidelines based on our own experiences, our research on responsible AI, and by listening to our customers and partners.
- As for facial recognition technology, developers using our Azure Face API can leverage [guidance](#) we've published to help them better understand its capabilities and limitations, ways to influence accuracy, and the importance of considering how the system will be deployed and used. We've also established [six principles](#) to govern the way we develop and deploy facial recognition at Microsoft, and we encourage other organizations to consider establishing guiding principles too. While companies all bear responsibility for exercising responsibility in this area, we have [called for increased legal regulation](#) and governance to broadly mitigate the risks of misuse.

These guidelines we have established will evolve over time as we at Microsoft continue to learn from our own experience and from the experiences of other tech companies, our customers, academics, civil society, and multi-stakeholder organizations as well. For more information and links to these guidelines, see the summary and resources section of this module.

## Engineering tools for responsible AI

Along with these guidelines, your governance system should take advantage of **tools and resources** that make it easier for developers to spot and design against potentially harmful issues like biases, safety and privacy gaps, and exclusionary practices. Below are some tools that we have found helpful.

### Fairness

AI systems should treat everyone fairly and avoid affecting similarly situated groups of people in different ways. There are two key steps for achieving this—assessment and mitigation:

#### Assessing fairness

1. Fairness in Machine Learning (ML) Systems ([FairLearn](#)) is an approach created by Microsoft Research and co-developed with products teams. FairLearn can be used to assess the potential unfairness of ML systems that make decisions about allocating resources, opportunities, or information. Note that fairness is a fundamentally sociotechnical challenge, so "fair" classification tools are not be-all-and-end-all solutions, and they are only appropriate in particular, limited, circumstances. A Python package that implements this approach is available on [GitHub](#).
  - a. For example, consider a ML system tasked with choosing applicants to interview for a job. FairLearn can turn a classifier that predicts who should be interviewed based on previous hiring decisions into a classifier that predicts who should be interviewed while also respecting demographic parity (or another fairness definition).
2. [Aequitas](#) is an open-source bias audit toolkit developed by the [Center for Data Science and Public Policy](#) at the University of Chicago for machine learning developers, analysts, and policymakers to audit machine learning models for discrimination and bias, and make informed and equitable decisions around developing and deploying predictive risk-assessment tools.



3. To understand the unique challenges regarding fairness in ML, watch our free [webinar on Machine Learning and Fairness](#). In this webinar you'll learn how to make detecting and mitigating biases a first-order priority in your development and deployment of ML systems.
4. For more on how organizations should approach assessing the fairness of their AI models, watch this NIPS [keynote address](#) from Kate Crawford, Principle Researcher at Microsoft and Co-founder of the AI Now Institute at NYU.

## Mitigating bias

1. A [methodology for reducing bias in word embedding](#) created by Microsoft Research helps reduce gender biases by modifying embeddings to remove gender stereotypes, such as the association between receptionist and female.
2. Read this paper from the ACM Conference on Fairness, Accountability, and Transparency: [Fairness and Abstraction in Sociotechnical Systems](#) which explains five key "traps" of fair-ML work and how to avoid them.
3. Read this paper from Cornell University: [Counterfactual Fairness](#) for an example of a framework for modeling fairness using tools from causal interference, and how it applies to the fair prediction of student success in law school.

## Reliability and safety

AI systems can become unreliable or inaccurate if their development and testing environment is not the same as the real world or if the system is not maintained properly. This can be especially dangerous in industries where safety may be at risk, like manufacturing or healthcare. To prevent reliability and safety issues, there are a number of technologies and tools that strengthen model performance through long-term monitoring and management:

1. The [Data Drift Monitoring](#) feature in [Azure Machine Learning](#) detects changes in the distribution of data that may cause degraded prediction performance, enabling developers to maintain accuracy by adapting the model to reflect changing data.
2. [Pandora](#) is a debugging framework designed by Microsoft Research to identify reliability and bias problems within machine learning models. It uses interpretable machine learning techniques, such as decision trees, to discover patterns and identify potential issues.
3. Microsoft [AirSim](#) is a valuable open-source tool for improving simulated training environments.

## Security and privacy

Security and privacy are key pillars of trust. To win the confidence of your customers and stakeholders, use the following resources to help protect security and privacy:

1. [Securing the Future of Artificial Intelligence and Machine Learning at Microsoft](#) provides guidance on how to protect algorithms, data, and services from new AI-specific security threats. While security is a constantly changing field, this paper outlines emerging engineering challenges and shares initial thoughts on potential remediation.
2. Secure execution environments such as [Azure confidential computing](#) help users secure data while it's "in use" on public cloud platforms (a state required for efficient processing). The data is protected inside a Trusted Execution Environment (TEE), also known as an enclave, such that code and data are protected against viewing and modification from outside of the TEE. This has a number of benefits, including the ability to train AI models using data sources from different organizations without sacrificing data confidentiality.
  - a. The Azure team has worked with Microsoft Research, Intel, Windows, and our Developer Tools group to develop our confidential computing solution, which enables developers to take advantage of different TEEs without having to change their code.
  - b. The Open Enclave SDK project provides a consistent API surface for developing apps using enclave-based computing.

3. Homomorphic encryption is a special type of encryption technique that allows users to compute on encrypted data without decrypting it. The results of the computations are encrypted and can be revealed only by the owner of the decryption key. To further the use of this important encryption technique, we developed [Microsoft SEAL](#) and made it open-source.
4. Multi-party computation (MPC) allows a set of parties to share encrypted data and algorithms with each other while preserving input privacy and ensuring that no party sees information about other members. For example, with MPC we can build a system that analyzes data from three different hospitals without any of them gaining access to each other's health data.
5. [Differential privacy](#) is a key technology for training machine learning models using private data. A differentially private algorithm uses random noise to ensure that the model output doesn't noticeably change when one individual in the dataset changes. This prevents attackers from inferring an individual's private information from the model's output.
  - a. [The PSI \(Private data Sharing Interface\) tool](#), developed by Harvard researchers, leverages differential privacy to enable researchers from many fields to explore and share datasets that contain private information.

## Inclusiveness

Inclusive design practices help ensure that AI models perform well for all users and no one is excluded from the opportunities provided by intelligent solutions. To help address potential barriers in your product environment that could unintentionally exclude people, use the following resources:

1. The Microsoft Research paper [Algorithmic Greenlining](#) proposes an approach that helps decision-makers develop selection criteria yielding high-quality and diverse results in contexts such as college admissions, hiring, and image search.
  - a. Take, for example, choosing job candidate search criteria. There's typically limited information about any candidate's "true quality." An employer's intuition might suggest searching for "computer programmer," which yields high-quality candidates but might return few female candidates. The algorithmic framework suggests alternative queries which are similar but more gender-diverse, such as "software developer" or "web developer."
2. Reference the Microsoft [inclusive design toolkit](#) and [inclusive design practices](#) to learn how to understand and address potential barriers in a product environment that could unintentionally exclude people.

## Transparency

The black-box nature of AI can be problematic and potentially harmful. To help your organization articulate how your AI models reach conclusions and build trust with your users, use the following resources:

1. [InterpretML](#) is an open-source package created by Microsoft Research for training interpretable models and explaining black box systems. It implements a number of intelligible models including Explainable Boosting Machine (EBM), an improvement over generalized additive models that has both high accuracy and intelligibility. It also supports several methods for generating explanations of black box model behavior or predictions including 'SHapley Additive exPlanations' (SHAP) and 'Local Interpretable Model-agnostic Explanations' (LIME).
2. [Azure Machine Learning](#) has a variety of tools that support model transparency. The [Model Interpretability](#) feature enables model designers and evaluators to explain why a model makes the predictions it does, which can be used to debug the model, validate that its behavior matches objectives, and check for bias.

## Accountability

The people who design and deploy AI systems must be accountable for how their systems operate. A useful first step for developing accountability and transparency in your organization is to create thorough documentation

processes for AI systems. To develop accountability practices for your own organization, leverage the following resources:

1. [Datasheets for datasets](#) is a paper that encourages people assembling training datasets to generate a datasheet with key information such as the motivation, composition, collection process, and recommended uses. Datasheets for datasets have the potential to increase transparency and accountability within the machine learning community, mitigate unwanted biases in machine learning systems, facilitate greater reproducibility of machine learning results, and help researchers and practitioners select more appropriate datasets for their chosen tasks.
2. The DevOps feature in [Azure Machine Learning](#) (called MLOps) makes it easier to track, reproduce, and share models and their version histories. It offers centralized management throughout the entire model development process, and helps teams monitor model performance by collecting application and model telemetry.
3. The Partnership on AI (PAI) is leading a multi-stakeholder initiative called [ABOUT ML](#) to develop, test, and promulgate best practices for machine learning documentation. These best practices may include documenting how AI systems were designed and for what purposes, where their data came from and why that data was chosen, how they were trained, tested, and corrected, and what purposes they're not suitable for.

## Engaging with external stakeholders

As the use of AI becomes more common, we consider it a shared responsibility across the public and private sectors to engage with AI responsibly. While each organization may have its own principles, governance systems, and guidelines to accomplish this internally, collaboration between enterprises, public organizations, governments, and non-profits will be crucial to ensure these concerns are addressed throughout society while maximizing AI's potential to deliver broad benefits. Organizations can contribute to these collective efforts in a number of ways. At Microsoft, we have focused on joining industry initiatives, influencing policy, addressing future labor and workplace needs, and considering how our technologies can be used to improve the lives of people around the world.

### Participating in industry coalitions

Advances in AI won't happen at just one company and they won't happen responsibly without a strong community working together. That's why enterprises should also take advantage of and contribute to shared initiatives like [Partnership on AI](#) (PAI). Partnership on AI is a group of researchers, non-profits, non-governmental organizations (NGOs), and companies dedicated to ensuring that AI is developed and utilized in a responsible manner. PAI is working to advance the public's understanding of AI, serve as an open platform for discussion and engagement about AI and its influences, and develop best practices for AI technologies. By sharing this knowledge with the world, organizations like PAI can help ensure that intelligent technologies are serving humanity in beneficial and responsible ways. As AI continues to expand, every enterprise should be open to collaboration and join in this shared dialogue.

### Providing a perspective on policy

Given AI is still at a nascent stage of development it is likely that changes to existing laws will be required and new laws may be needed to keep AI on a path that benefits society. We believe it's important for organizations to be aware of these potential changes and get involved where possible in the public dialogue on AI advancements and potential legislation.



For example, last year we publicly [called for regulation](#) of facial recognition technology and outlined our recommendations for the public and private sector alike. Over the past decade, facial recognition technology has made rapid advancements in both speed and accuracy, uncovering myriad uses ranging from cataloging photos to enhancing security. Its potential for uses that infringe on fundamental human rights like open public surveillance has highlighted the need for societal scrutiny and government regulation. We believe there is an especially pressing need for laws and regulations around AI technology with sensitive use cases like this. Other emerging intelligent technologies will raise similar issues in the future, and we will continue to look for opportunities to help government agencies create adequate regulations. We will all need to continue to work together to identify issues that have clear societal or economic consequences and develop solutions to protect people without restricting innovation.

In addition, we should acknowledge the broad concerns about the impact of AI technologies on jobs and take steps to ensure people are prepared for the impact of these technologies in the workplace and the workforce.

## Addressing future labor and workplace needs

AI is already significantly shifting how people prepare for, find, and accomplish work. Increasingly, workers will have to learn how to maximize the benefits of using emerging intelligent technology in their workplace. Gartner predicts that by 2021, 70% of enterprises will integrate AI to assist employee's productivity.<sup>ii</sup> But like other major new technologies, AI raises concerns about its long-term impact on employment. A [recent study](#) by Accenture concluded that "AI could double annual economic growth rates in 2035 by changing the nature of work and creating a new relationship between man and machine."<sup>iii</sup> That's why, at Microsoft, we are focused on creating AI that augments human abilities rather than replacing them. We aim to empower the workforce by developing AI technology that leverages the unique strengths of computers—such as probabilistic reasoning and pattern recognition—with the creativity, ingenuity, and capacity for meaning-making of humans.

Yet as the pace of technological innovation and disruption is accelerating, we are seeing a widening skills gap in critical areas of foundational knowledge and expertise. The McKinsey Global Institute estimates that, by 2024, there will be 250,000 unfilled U.S. data science job openings.<sup>ii</sup> To realize the potential of AI, we must first address the ability of those in the current and future workforce to adapt to, engage with, and fundamentally understand these technological advances and transformations.

We believe organizations have an important role to play in helping people get the training they need to thrive in today's shifting economy. At Microsoft, we are focusing on three areas:

1. Preparing students for tomorrow's jobs
  - a. Our [Technology Education and Literacy in Schools \(TEALS\)](#) program recruits volunteers from across the tech industry to team-teach with classroom teachers, helping address the widespread shortage of computer science teachers. Since its founding in 2009, TEALS has expanded across the U.S. and served more than 37,000 students.
2. Helping workers prepare for the changing economy
  - a. Providing publicly available resources aimed at helping workers increase their AI-specific skills. For example, our [Microsoft Professional Program](#) now has an AI track bringing together expert instructors, provide hands-on labs, offer AI-specific online courses and instructional videos. We have also created a developer-focused [AI School](#), which provides online videos and other assets that help build professional AI skills.
3. Helping match workers to 21st century job opportunities.
  - a. We are partnering with The Markle Foundation on the [Skillful Initiative](#). This program helps workers identify high-demand jobs and the skills they need to fill them, helps employers find the talent they

need to grow, and helps educators train people with the skills required to participate in today's economy.

We recognize that these programs and investments are just the beginning. As long as technological advancement continues, so will the possibility of disruption. If we want to build a future with an inclusive view of opportunity, well-being, and broadly-available work with dignity and purpose, we need to build greater resiliency to transformation by fostering structural supports for lifelong learning.

## Contributing solutions to societal challenges

By harnessing the power of AI, we have an opportunity to help individuals, organizations, and governments solve some of society's most daunting problems. AI can help people collect, process, and analyze data, turning it into actionable insights that can accelerate advances in health and wellbeing, education, agriculture, and climate change mitigation. We invite all enterprises to consider making their AI-infused technologies widely available to improve the lives of people around the world. At Microsoft, we are providing technology resources and expertise to empower those working to preserve cultural heritage, solve humanitarian issues, improve the health of communities, and, create a more sustainable and accessible world through our AI for Good initiative. We've committed \$165 million over the next five years to provide financial grants, technology investments, and partnerships that combine our AI and data science experience with these individuals' and organizations' core expertise. AI for Good includes five programs that are already supporting over 600 projects across the globe: AI for Accessibility, AI for Earth, AI for Humanitarian Action, AI for Cultural Heritage, and AI for Health.

### AI for Accessibility

AI can be a game-changer for people with disabilities. Intelligent solutions that enable people to leverage computers to see, hear, and reason are already impacting people's lives around the world. Currently, however, only one in ten people with disabilities have access to assistive technologies and solutions. Driven by our own principle of inclusiveness, we created [AI for Accessibility](#), a commitment to placing AI tools in the hands of developers working to create accessible solutions for the 1 billion-plus people with disabilities around the world.

Over the last two decades, we have been working to foster communication and break down language barriers with [Microsoft Translator](#). To help accurately translate more than 60 languages, it uses an advanced form of automatic speech recognition to convert raw spoken language into fluent punctuated text. To further explore this promising application, we partnered with the National Technical Institute for the Deaf (NTID) in New York to pilot the use of Microsoft Translator in classrooms. Combined with the American Sign Language translators, the captions provide another layer of communication to help the 1,500 students attending NTID learn. NTID students have also started using the Microsoft Translator app outside of class to help them communicate with hearing peers.

### AI for Earth

We face a collective need for urgent action to address global climate issues. From conserving water to protecting biodiversity, we believe AI can help us better understand and address environmental issues. Our [AI for Earth](#) initiative has committed to provide AI to those who are working to solve environmental issues in four key areas: climate, water, agriculture, and biodiversity.

One such project is called [FarmBeats](#). We are working with farmers around the world to provide them with AI that can help them monitor the health of their farms in real time, enabling data-driven decisions that can help increase agricultural yield, lower costs, and reduce environmental impact. Even on farms without power or internet access, the FarmBeats program can use drones and low-cost sensors to capture data. The FarmBeats team can then apply machine learning algorithms to integrate sensor data with aerial imagery from drones and other relevant data

(such as weather, crop predictions, and best practices) to deliver actionable insights to farmers at a fraction of the cost of existing solutions. The Dancing Crow Farm in Carnation, Washington, has already lowered water input and increased yield using this program. With the world's population expected to grow by nearly 2.5 billion people over the next quarter century, AI solutions like FarmBeats offers significant opportunities to increase food production by improving agricultural yield and reducing waste.

## **AI for Humanitarian Action**

With intelligent technology, we have a tremendous opportunity to impact how organizations anticipate, predict, and respond to pressing humanitarian crises. That's why we are dedicating our [AI for Humanitarian Action](#) initiative to expediting new AI solutions in four areas: disaster recovery, needs of children, refugees and displaced people, and human rights. By empowering global relief organizations with intelligent technology, we can help save more lives, alleviate suffering, and restore human dignity.

For more than six years, Microsoft has been partnering with [Operation Smile](#) to help them deliver the best possible care. Since its founding in 1982, the global nonprofit has provided free and safe surgeries in low- and middle-income countries for more than 270,000 children and young adults with cleft lip, cleft palate, and other facial deformities. Currently, Operation Smile uses a solution with SharePoint and Power BI that has streamlined workflows across the company and cut post-surgery evaluation time in half, from four to two months.

## **AI for Cultural Heritage**

Technology has played a big role in accelerating globalization. While it's our business to advance technology, we also believe that technology should respect and even help protect the world's timeless values. That conviction has led us to add [AI for Cultural Heritage](#) to our AI for Good portfolio.

Our AI for Cultural Heritage program will incentivize the use of artificial intelligence by nonprofits, universities, and governments around the world to help preserve the languages we speak, the places we live, and the artifacts we treasure. It will build on recent work we've pursued using various aspects of AI in each of these areas. For example, we have collaborated with The Metropolitan Museum of Art and MIT to explore ways in which AI can make The Met's Open Access collection accessible, discoverable, and useful to the 3.9 billion internet-connected people worldwide.

## **AI for Health**

Health is a global issue that impacts every person and transcends every border. Technology plays an important role in advancing research and improving access to care for underserved populations. We are aiming to tackle some of the toughest challenges in health by collaborating with nonprofits and researchers through our [AI for Health](#) initiative.

The initiative will focus on accelerating the quest for discovery of potential treatments, increasing our shared understanding of health and longevity, and reducing health inequity. The program is underpinned with a strong foundation of privacy, security and ethics, and was developed in collaboration with leading health experts who are driving important medical initiatives. We have partnered with clinicians, researchers, and health professionals to make the value of AI more accessible and impactful. For example, we're partnering with the Novartis Foundation to teach an AI model to recognize leprosy in an image of a skin lesion. Similarly, Intelligent Retinal Imaging Systems (IRIS) is developing an AI solution to autonomously diagnose diabetic retinopathy by evaluating images.





## Working together for a better future

From holistically transforming industries to addressing critical issues facing humanity, AI is already solving some of our most complex challenges and redefining how humans and technology interact.

In this paper, we walked through some of the steps Microsoft is taking to prioritize responsible AI in the hope that our experience can help others. However, we recognize that we do not have all the answers and every individual, company, and region will have their own unique beliefs and standards that should be reflected in their path towards responsible AI. The structures, processes, tools, and resources mentioned in this paper can be a starting point from which organizations can adapt to create their own AI strategy.

Early adopters of AI Organizations that embrace AI early have an important role to play in promoting the responsible use of AI and preparing society for its impacts. Their firsthand experience in dealing with the ethical challenges of AI will be crucial knowledge for later adopters and those trying to study or regulate AI technology.

As organizations and as a society, our steps towards responsible AI will need to continually evolve to reflect new innovations and lessons from our mistakes and accomplishments. It's essential that we continue to foster open dialogue among businesses, governments, NGOs, academic researchers, and all other interested individuals and organizations.

We look forward to working collaboratively to develop and share methods to ensure AI will benefit society. By engaging with AI in a responsible manner, we can ensure that it fulfills its promise to create a better future for all.

## Learn more today

To learn more about how to take actionable steps toward responsible AI, please visit our [AI Business School](#).

The AI Business School is a master class series designed specifically for business leaders that will empower them to be successful and get results from AI. The online class series is free, on-demand, and includes content built by business professors from INSEAD, business leaders from Microsoft and across industries, and Microsoft technology experts. It represents their latest thinking on AI transformation in the enterprise and how to set your organization up for success.

The AI Business School has four modules that will help you understand the diverse array of AI technology, define an AI strategy to create business value, discover ways to enable an AI-ready culture in your business, and give you actionable steps to help you foster responsible and trustworthy AI.

For more general information on this subject, we invite readers to explore the documents and resources below.

## ***The Future Computed: Artificial Intelligence and Its Role in Society***

A short non-technical book discussing AI's promise and challenge for modern society, by Microsoft President Brad Smith and top AI researcher Harry Shum.

<https://blogs.microsoft.com/blog/2018/01/17/future-computed-artificial-intelligence-role-society/>

## ***AI Business School***

An engaging master class series designed specifically to help business leaders address key challenges around developing and implementing an AI-driven approach

<https://www.microsoft.com/en-us/ai/ai-business-school>

## ***Facial Recognition: It's Time for Action***

Microsoft President Brad Smith makes the case for careful regulation of AI-based facial recognition and an ethical approach by governments and corporations to all AI technologies.

<https://blogs.microsoft.com/on-the-issues/2018/12/06/facial-recognition-its-time-for-action/>

## ***Digital Transformation in the Cloud: What enterprise leaders and their legal and compliance advisors need to know***

A non-technical book that lays out the business case for digital transformation and explains the foundational requirements of security, privacy, and compliance.

<https://aka.ms/digital-transformation-in-the-cloud/>

## ***Microsoft Vision for AI in the Enterprise***

A whitepaper that outlines how enterprises can bring AI to every application, business process, and employee.

<https://www.microsoft.com/en-us/ai/business>

## ***A Cloud for Global Good***

Microsoft's policy considerations and recommendations for creating a framework of laws that extend the benefits of the cloud to all.

<https://news.microsoft.com/cloudforgood/>

## ***Transformation Tuesday blog series***

A series of blog posts by Michael McLoughlin—Microsoft on digital transformation, AI, and cloud computing.

<https://aka.ms/Transformation-Tuesday>

## ***Partnership on AI***

The Partnership on AI was established to study and formulate best practices on AI technologies, to advance the public's understanding of AI, and to serve as an open platform for discussion and engagement about AI and its influences on people and society.

<https://www.partnershiponai.org/>

## ***Microsoft AI***

A more technical introduction to Microsoft AI.

<https://www.microsoft.com/en-us/ai>

## ***The AI blog***

The latest news about AI progress from Microsoft.

<https://blogs.microsoft.com/ai/>

## ***AI Insight Series: A Short Introduction to Artificial Intelligence***

An introduction to AI for legal and compliance professionals.

<https://aka.ms/short-intro-to-ai/>

## ***AI Insight Series: Artificial Intelligence and the GDPR Challenge***

Considerations to deploying AI solutions within the framework of EU's General Data Protection Regulation.

<https://digitaltransformation.instantmagazine.com/pub/ai-and-gdpr>

## ***Artificial Intelligence—The ethical and Legal Implications***

A portal with articles and resources for legal and compliance professionals by Norton Rose Fulbright related to ethics and AI.

<https://www.insidetechlaw.com/artificial-intelligence/>

## ***IAPP webcast: Artificial Intelligence, Machine Learning, and Data Ethics***

Hear three experts with legal, technological and ethical perspectives wrestle with the challenge of Artificial Intelligence, machine learning, and data ethics issues.

<https://iapp.org/store/webconferences/a011P00000CVdPnQAL>

---

<sup>i</sup> Inc., “31 Tech Predictions for 2019.” Christina DesMarais, 12 December 2018. <https://www.inc.com/christina-desmarais/31-tech-predictions-for-2019.html> and Forbes, “Five Trends Shaping the Future of Customer Experience in 2019.” Blake Morgan, 5 December 2018. <https://www.forbes.com/sites/blakemorgan/2018/12/05/five-trends-shaping-the-future-of-customer-experience-in-2019/#754a5cab7bb6>

<sup>ii</sup> Gartner Press Release, “Gartner Predicts 70 Percent of Organizations Will Integrate AI to Assist Employees’ Productivity by 2021,” January 24, 2019. <https://www.gartner.com/en/newsroom/press-releases/2019-01-24-gartner-predicts-70-percent-of-organizations-will-int>

<sup>iii</sup> Accenture, “Artificial Intelligence is the Future of Growth.” Mark Purdy and Paul Daugherty, September 2016. <https://www.accenture.com/us-en/insight-artificial-intelligence-future-growth>

<sup>ii</sup> McKinsey Global Institute, “The age of analytics: Competing in a data-driven world.” Nicolaus Henke, Jacques Bughin, Michael Chui, James Manyika, Tamim Saleh, Bill Wiseman, and Guru Sethupathy, December 2016. <https://www.mckinsey.com/business-functions/mckinsey-analytics/our-insights/the-age-of-analytics-competing-in-a-data-driven-world>



© 2020 Microsoft Corporation. All rights reserved. This whitepaper is provided “as-is.” Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it.

Some examples are for illustration only and are fictitious. No real association is intended or inferred.

This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.