IBM Office of Privacy and Responsible Technology

Reflecting on the Five-Year Anniversary of IBM's AI Ethics Board

Table of contents

Introduction

In the years since IBM first formed its AI Ethics Board in 2019 as a centerpiece of our AI governance process, our governance journey has continuously evolved as emerging technologies like artificial intelligence (AI) have become ubiquitous across industries.

Our progress in this space started years before that and has been marked by key moments of leadership ever since. IBM appointed one of the industry's first AI Ethics Global Leaders in 2015, co-founded the Partnership on AI in 2016, announced our Principles for Trust and Transparency in 2018, published our internal Ethics by Design playbook in 2021, and co-launched the AI Alliance in 2023. These events reflect some of the many steps IBM has taken to maintain its commitment to being a responsible steward of technology. We build systems based on trust by embedding ethical principles into data and AI across the development, deployment, and use of technology.

As society moves from cautious exploration to rapid adoption of AI at scale, trust and transparency are more important than ever. IBM maintains its commitment to these principles through organizational and technical governance of AI.

The work of the IBM AI Ethics Board underscores the alignment between company values and ethics, which are integrated into our policies, products, and strategy through a diverse and multidisciplinary AI Ethics Board membership. Serving as Co-Chairs of the Board is a great honor and a meaningful experience.

Christina Montgomery
Chief Privacy & Trust Officer
Francesca Rossi
AI Ethics Global Leader

AI Ethics Board Co-Chairs

IBM’s AI Governance

Our company’s AI strategy is supported by organizational governance – shaping a culture of trustworthy AI, establishing principles, applying guardrails, and enabling regulatory compliance, as well as technical governance – implementing tools and processes to manage and monitor AI governance across the AI lifecycle.

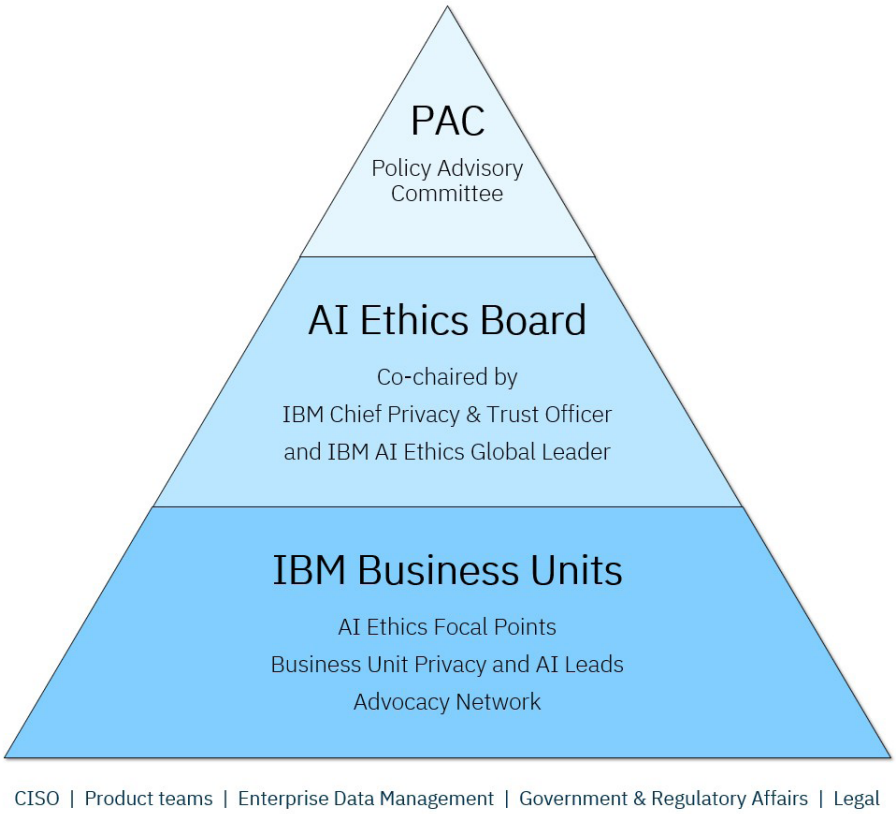

IBM's multi-layered organizational AI governance structure helps infuse AI ethics throughout the company

The IBM AI Ethics Board is one critical component of our organizational governance, and IBM technology enables our AI lifecycle management to consistently apply, evaluate, monitor, and track our AI governance and compliance workstreams.

The IBM AI Ethics Board mission is to provide governance and decision-making as IBM develops, deploys and uses AI (and other technologies) to maintain consistency with the company's values and to advance trustworthy AI for our clients, partners, and the world. The Board is overseen by the Policy Advisory Committee (PAC), a group of IBM's most senior leaders.

The Board is supported by various groups and teams, including AI Ethics Focal Points, Privacy and AI Leads in each business unit, and a grassroots Advocacy Network of subject matter experts from across the company, including members from our Business Resources Groups (BRGs). The BRGs bring a diverse perspective, as they are employee-led groups formed around a common interest, bond, or background that help inform the Board's positions and points of view. The Board and the AI Ethics Focal Points work closely together in a top-down and bottom-up approach, and are the two groups most responsible for reviewing the potential ethical risks of proposed AI use cases.

> The partnership between the Board and the AI Ethics Focal Points is essential. By working together, we make certain that every AI use case not only reflects our core values but is also meticulously assessed for ethical implications, promoting a culture of responsible and innovative progress.
>
> **Kitty Chaney Reed**
> Chief Leadership, Culture and Inclusion Officer
> IBM AI Ethics Board Member

The AI Ethics Board inside IBM’s Integrated Governance Program

We have expanded our existing privacy program to build our Integrated Governance Program (IGP), allowing us to adapt quickly to address new AI requirements and obligations by using many existing processes and systems.

IBM's governance framework coupled with our strong technology foundation enable us to build and adopt responsible AI at scale across the company

IGP has allowed IBM to move from a more reactive to a continuous compliance model for data, privacy, and AI. This continuous compliance approach empowers us to bring AI products and services to market with speed and trust, for ourselves and for our clients. IGP is underpinned by IBM's own technology including watsonx, serving as a living lab for our trusted solutions. Integrating into a single program with a high degree of accountability within business units helps enable us to keep pace with future requirements while serving the current needs of the business.

Companies need to consider an integrated governance approach as the foundation of their AI strategy. An Integrated Governance Program can enable companies to develop, deploy and maintain AI systems in a way that aligns with their values, risk tolerance and regulatory requirements. A diverse and multidisciplinary AI Ethics Board can play a critical role in an Integrated Governance Program by helping to ensure AI use cases align with a company's principles and values.

Christina Montgomery
Chief Privacy & Trust Officer
AI Ethics Board Co-Chair

The Influence of the AI Ethics Board

The work of the Board drives IBM to continue its leadership in AI and tech ethics, by innovating responsibly across four areas: in the marketplace, informing policies and regulations, conducting world-class research, and developing points of view on emerging technologies.

Business

One of the primary workstreams of the Board is the Tech Ethics Use Case Assessment. This process enables a diverse body of stakeholders to review AI and other technology use cases and implement guardrails to drive alignment with our AI ethics principles. AI Ethics Focal Points act as the first point of contact within their business units to proactively identify and assess AI ethics concerns, mitigate risks for individual use cases, and forward projects as necessary to the IBM AI Ethics Board for further review. This assessment evaluates aspects like the type of data used to develop the technology, where the technology will be used and by whom, and the potential for unintended secondary uses of the technology that could cause harm. Based on this assessment, appropriate guardrails are established for the use case.

The Board is an integral part of IBM's AI governance process, including contributing to an AI Baseline that identifies common AI principles from regulations and voluntary commitments. The Baseline is part of the Integrated Governance Program that provides organizational and technical governance of AI across the enterprise.

To instill and maintain a culture of ethics, the Board supports IBM's AI ethics education program which creates learning experiences that are tailored to different audiences to help them understand how AI ethics applies to their work. Internally, all IBM employees receive AI ethics education through IBM's required annual education and can choose to take a deeper dive by earning a foundations-level badge. Beyond our internal education, IBM educates and trains partners on AI ethics as part of Trust & Compliance Integrity Courses and Integrity Education, and we are making excellent progress on the goal we announced in IBM's Impact Report to train 1,000 technology suppliers in technology ethics by 2025, with more than 600 suppliers trained at the end of 2023.

Under our IBM Impact framework, we are advancing this goal beyond our suppliers. Executives estimate that implementing AI and automation will require companies to reskill 40% of the workforce over the next three years. To help address this challenge, IBM committed to training two million learners in AI globally by 2026. Through channels like our free education program, IBM SkillsBuild, we offer courses on topics including AI ethics, providing real-world examples that guide learners through the five pillars of trust: fairness, robustness, explainability, transparency, and privacy.

> IBM has strived for more than a century to introduce new technologies like AI responsibly and with a purpose. We are pioneers in our industry, establishing an AI Ethics Board and infusing IBM's principles into our products to usher these powerful innovations into the world in an ethical way. Beyond critical internal governance practices, one important component in every AI journey is to prepare the workforce. The talent gap is one of the biggest challenges facing businesses today, and it will take bold and quick action to address that need, including skilling millions of people so that they embrace and understand AI and AI ethics.
>
> **Justina Nixon-Saintil**
> Corporate Social Responsibility and Chief Impact Officer
> IBM AI Ethics Board Member

I'm proud to serve on IBM's AI Ethics Board, which helps build public trust by promoting responsible and purposeful deployment of new technologies. This cross-disciplinary group has fostered a culture of ethical tech within IBM and has helped to advance smart public policy. The Board engages with diverse policy voices, and educates lawmakers on AI regulation.

Roslyn Docktor
Technology and Science Policy
IBM AI Ethics Board Member

Policy

We not only apply our principles throughout our business, but also advocate for policies that promote AI ethics globally. We are actively working with regulators and industry leaders to shape AI governance, regulation, and standards. For years, we have advocated for a precision regulation approach to AI, which our CEO expanded upon in recommendations for advancing trusted AI: regulating AI risk and not AI algorithms; making AI creators and deployers accountable, not immune to liability; and supporting open AI innovation, not an AI licensing regime. In May 2023, our Chief Privacy and Trust Officer Christina Montgomery brought this stance to the US Senate Judiciary Subcommittee, testifying about the manifest benefits of AI that are possible when precision regulation is applied.

We have also taken points of view on the opportunities, risks, and mitigations of foundation models as well as what policymakers should consider when it comes to foundation models, offering our expertise and thought leadership for consideration as the industry and world alike adapt to the era of generative AI.

At IBM, we believe in an open innovation ecosystem for AI. An open innovation ecosystem is one in which stakeholders recognize the value of community-built technology and the open exchange of information, ideas, and skills it cultivates. This is also one of the goals of the AI Alliance. Open innovation doesn't just mean open-source software. Open source and permissively licensed AI models are a key part of an open innovation ecosystem for AI, as are open-source toolkits and resources, open datasets, open standards, and open science. Since 2018, IBM has been releasing resources to the open source community, including the AI Fairness 360, AI Explainability 360, and Adversarial Robustness 360 toolkits that were donated to the Linux Foundation in 2020. IBM believes open is safe, open is innovation, and open is opportunity.

Research

IBM Research serves as the company's organic growth engine and has played an integral role in advancing AI ethics. IBM Research has incubated several initiatives, including the Data and AI for Social Impact (DAISI) program, which empowers social change organizations to use AI to work toward their missions. IBM Research's initiatives, partnerships and products take a principled approach that begins with the publication of hundreds of peer-reviewed fundamental research papers. The insight and innovation gained from these research initiatives and partnerships is then infused into our product development, creating tooling to enable ethics and governance as in our watsonx platform. An example of this is the embedding of governance into the data and model factory that produced IBM's Granite models.

Led by the Office of Privacy and Responsible Technology, IBM's collaboration with the Notre Dame-IBM Tech Ethics Lab is focused on applied research and best practices in technology ethics. The lab, a first-of-a-kind exclusively focused on the ethical implications of technology, promotes human values in technology through tangible, applied, and interdisciplinary research that addresses core ethical questions. Last year, IBM and Notre Dame published a paper on "The Return on Investment in AI Ethics: A Holistic Framework", which was recently highlighted in the California Management Review. The Tech Ethics Lab has also sponsored the Pulitzer Center's AI Spotlight Series to build AI literacy in ways that are relevant and accessible to a non-technical journalism audience, with the goal of training 1,000 journalists over two years.

> The Research division takes on the difficult role of advancing the theory of AI ethics, while also figuring out the tools and processes for operationalizing it through trustworthy AI and AI governance.
>
> **Kush Varshney**
> IBM Fellow
> IBM AI Ethics Board Member

As generative AI reshapes industries, it's imperative that we prioritize responsible development and governance. At IBM, we are committed to addressing the complexities of AI risks, enabling our solutions to empower businesses to innovate confidently and ethically. The insights from our governance products like watsonx.governance are designed to guide organizations in navigating the evolving AI landscape, turning potential challenges into opportunities for sustainable growth.

Ritika Gunnar
General Manager Product Management for Data and AI Software
IBM AI Ethics Board Member

Insights

The rise of generative AI in late 2022 brought AI to the cultural forefront. Leaders across all industries saw the promise of generative AI-enabled business transformation. It was clear from the outset that as generative AI evolved, so too did the AI risk landscape. These potential risks have significant implications for how AI is designed, developed, deployed, used, and governed. The IBM AI Ethics Board's white paper "Foundation models: Opportunities, risks and mitigations" describes the new, amplified, and traditional risks embodied by generative AI as well as risk mitigation strategies to consider. IBM's watsonx AI Risk Atlas, based on this white paper, identifies over 60 unique risks of working with generative AI, foundation models, and machine learning models across all phases of the AI lifecycle. The Risk Atlas content is now also available in watsonx.governance, with a Risk Identification Assessment questionnaire that identifies potential risks applicable to the use case, creating a risk profile to enable the implementation of mitigations and guardrails. This is a prime example of the AI Ethics Board's work contributing back to IBM products and research.

Putting principles into practice, the IBM AI Ethics Board also published a white paper "Augmenting Human Intelligence - the IBM Point of View" operationalizing IBM's first Principle for Trust and Transparency, that the purpose of AI is to augment human intelligence. AI should be designed to include and balance human oversight, agency and accountability over decisions across the AI lifecycle. In addition to the white paper's case studies, best practices and Key Performance Indicators for augmenting human intelligence with AI, the Board has published examples of KPIs in action, and a deep dive into IBM's Consulting Assistants and how these supplement the capabilities, experiences, and insights of humans.

Lessons Learned and Best Practices

As AI and technology continue to evolve, IBM's AI governance journey has evolved as well.

Key best practices that IBM has developed over the past five years include:

- Embedding AI ethics across the company through business unit AI Ethics Focal Points and a grassroots Tech Ethics Advocacy Network.
- Instilling a culture of ethics through role-based education programs for different audiences with varying levels of familiarity with AI, both internally and externally.
- Strengthening alignment with AI ethics principles and AI governance processes by providing assessments of the risks of AI solutions, as well as mitigation strategies and tools.
- Integrating governance methodologies and tools, and using our own technology to automate where appropriate, for the implementation and monitoring of AI governance throughout the development, deployment, and use of trustworthy AI.
- Leveraging and adapting existing governance mechanisms (like privacy programs) to help address AI regulatory requirements and enable continuous compliance.

As we mark the fifth anniversary of the IBM AI Ethics Board, we've taken a few moments to reflect on the progress the Board has made to keep the trust of our clients. Over the years we've engaged in spirited discussions, explored best practices and collaborated with governments and organizations across the globe. Together, our collective efforts have helped shape a culture of developing AI technology and solutions that encourage innovation while providing the protections that society needs.

There is increased scrutiny on both the opportunities and risks associated with AI, and it is in the best interest of our ecosystem to institute intentionality in how we consider its usage. As IBM's Chief Sustainability Officer, I am prioritizing our efforts in the intersection of AI and sustainability – how the technology can help support accelerating our climate and sustainability goals, and how we also need to consider the sustainability of the AI lifecycle. I am proud of IBM's leadership in this space, and believe it's our collective responsibility to continue to build trust and transparency with our partners. The AI Ethics Board has – and will continue to be – a critical partner in this ever-evolving journey.

Christina Shim
Chief Sustainability Officer
IBM AI Ethics Board Member

Resources

Learn more about IBM's approach to AI ethics:

- AI Ethics at IBM
- IBM AI Governance Consulting
- watsonx.governance
- 2023 IBM Impact Report
- Case study: Building trust in AI
- The CEO's Guide to Generative AI: Responsible AI and Ethics

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the
United States of America
November 2024