Collibra

Whitepaper

AI governance: An essential element for successful AI

June 2024

Table of contents

Why AI governance is essential

The AI era is here

Today, we're in the early stages of a fundamental economic, social and political transformation driven by revolutionary technological advances in artificial intelligence (AI). The launch of OpenAI's ChatGPT in November 2022 and its explosive growth marked a significant milestone in mainstream AI adoption. Since that time, the phenomenal success of AI hardware leader Nvidia, which has increased its market capitalization by more than $2T during this period, represents an additional indication of the reach and influence of the AI revolution. We are truly entering the AI era — and every forward-looking enterprise in the world is seeking ways to effectively integrate AI into their business.

> **64% of businesses expect AI to increase productivity**

The AI reality: Tremendous potential, real challenges

While the promise of generative AI gets all the spotlight, the truth is AI also presents significant challenges — including regulatory, data, organizational and ownership complexities — that necessitate a robust AI governance solution.

Companies that implement successful AI initiatives understand that a reliable AI governance framework is essential to:

- Improve compliance
- Reduce risk
- Optimize AI development
- Ensure the reliable, traceable and compliant use of AI

For data scientists, AI governance delivers these three key benefits:

- Eliminates the AI "black box," providing transparency into AI models and enabling model cataloging and performance tracking, aiding in efficient management and issue resolution
- Provides data leaders with a solution for ensuring compliance in a rapidly evolving regulatory environment
- Ensures data scientists have access to high-quality, trusted data

Most importantly, an AI governance framework provides a reliable, repeatable approach to AI initiatives.

AI governance, defined

AI governance is the application of rules, processes, and responsibilities to drive maximum value from your automated data products by ensuring applicable, streamlined, and ethical AI practices that mitigate risk, adhere to legal requirements, and protect privacy.

AI governance: Essential for AI use case development

Given the transformative impact and potential of AI, and the evolving regulatory environment, AI governance is an essential pillar of any AI roadmap. For enterprises pursuing the development of AI, AI governance has become a de facto standard for maximizing AI use case effectiveness and minimizing AI data risk.

[1] Source: https://www.forbes.com/advisor/business/software/ai-in-business

Challenges: The risks are real

It's clear that AI has the potential to transform industries, offering unparalleled opportunities for innovation and efficiency. However, the journey to realizing these benefits is fraught with challenges, such as:

1. Rapidly evolving regulatory environment
2. Poor data
3. Lack of organizational alignment
4. Challenges choosing an AI ownership model

1. Regulations: The regulatory environment is rapidly evolving

Corporate enterprises that want to use AI in the EU and the US must navigate a complex landscape of regulations that govern the ethical, transparent and responsible deployment of AI technologies.

In late May 2024, the Artificial Intelligence Act was approved by the EU, establishing a common AI regulatory and legal framework. The AI Act stipulates that enterprises that want to do business in the EU must ensure that AI systems, especially generative AI applications, meet the essential risk mitigation criteria.

Additionally, 2018's General Data Protection Regulation (GDPR) applies to all entities processing the personal data of EU citizens, and requires that organizations implement data protection principles. AI systems that process personal data must comply with GDPR, ensuring data is handled lawfully and ethically.

In the US, while the federal government has yet to issue a broad regulatory framework, the federal Office of Management and Budget (OMB) recently issued M-24-10, a memorandum that compels federal agencies to swiftly enact a number of specific AI management protocols, including appointing a Chief Data and AI Officer (CDAIO) and establishing an AI Governance Board.

And in May 2024, Colorado was the first state to enact consumer protection legislation that adopts a risk-based approach to AI (similar to the EU AI Act), requiring businesses that deploy "high-risk" AI systems to submit impact assessments and protective measures to prevent algorithmic discrimination.

Data quality is the primary barrier to AI adoption.

2. Poor data: Garbage (data) in, garbage (AI) out

In today's data-driven world, the daily volume of data that's generated around the world is staggering. By 2025, it's estimated that 463 exabytes of data will be created every day. But this overwhelming influx of data means that it is crucial for organizations to develop robust data management and governance strategies.

IT leaders—from the CDAO to the CTO, CIO, CISO to the CAIO (Chief AI Officer)—understand AI is built on data. And like any other data-powered product, data hygiene is key to everything. Indeed, access to AI-ready data remains one of the most significant hurdles in implementing AI solutions.

In fact, poor data quality can lead to a range of challenges for AI projects, including:

- Biased decision-making
- Hallucinations
- Inaccurate recommendations
- Security risks

For AI initiatives to succeed, organizations must prioritize improving data quality, ensuring that data is reliable, consistent and complete. At Collibra, we understand that data is at the heart of any data product, and nowhere is this more true than with AI models.

Ready to get started? Check out our AI Readiness Checklist.

What is a data product?

A data product is an asset that leverages data to generate actionable insights, drive decision making or deliver specific outcomes for users. It involves collecting and integrating data from various sources, processing and analyzing it to extract meaningful insights, and ensuring data quality and governance.

[1] Source: Forrester, Feb 2024

3. Organizational alignment: Getting on the same page about AI

Implementing AI use cases within an organization is not merely a technical endeavor — it involves significant cultural and procedural changes. The organizational challenge is often the most formidable barrier to AI adoption, as it requires shifting mindsets, updating processes and fostering collaboration across diverse teams.

Change is inherently difficult, and resistance can come from various levels within an organization. Employees may fear that AI will replace their jobs, or they may be skeptical about the benefits of new technologies. This resistance can manifest in several ways, including

- Reluctance to adopt new tools
- Unwillingness to participate in AI projects
- A general lack of enthusiasm for change

Effective change management is vital to overcome resistance and foster a culture that embraces AI. This involves breaking down silos, encouraging cross-functional collaboration and ensuring that all stakeholders understand the value and implications of AI use cases.

Strategies to foster collaboration

- Cross-functional teams: Establish cross-functional teams that include members from various departments. These teams can work together on AI projects, ensuring that diverse perspectives are considered and integrated
- Unified data strategy: Develop a unified data strategy that aligns with the organization's AI goals. This strategy should facilitate data sharing and integration across departments while maintaining data governance standards
- Regular communication: Encourage regular communication and collaboration between teams. This can be achieved through joint workshops, meetings and collaborative platforms that enable seamless information sharing

4. AI governance ownership: Data-, model- or compliance-centricity?

One of the biggest challenges organizations pursuing AI initiatives face is choosing and establishing who owns AI governance. AI governance can fall within three types:

- Data-centric
- Model-centric
- Compliance-centric

Data-centricity: Prioritizing quality and accessibility

A data-centric approach focuses on helping data, AI, security, risk and legal, and business teams deliver trusted AI by providing easy access to reliable data and implementing appropriate controls across the AI use-case lifecycle. This model demands close collaboration between teams to ensure that AI use cases are safely deployed into production. By working together, these teams can address potential risks, enforce compliance and enhance the overall quality of AI outputs.

The primary goal of this approach emphasizes the importance of high-quality data and ensuring effective AI use-case development and risk mitigation.

Model-centricity: Optimizing AI models

A model-centric approach helps AI and data teams implement AI use cases effectively by preparing, developing, running and monitoring AI models. This approach focuses on the technical aspects of AI, ensuring that models are accurate, reliable and efficient. In a model-centric framework, collaboration primarily occurs between AI and data teams, who work together to optimize model performance and address any technical challenges that arise during development and deployment.

The goal of a model-centric approach is to support the development and running of AI models.

Compliance-centricity: Ensuring legal and ethical integrity

A compliance-centric approach focuses on helping legal and privacy teams ensure adherence to laws and regulations by fully documenting and auditing the use of AI. In this framework, legal teams take the lead, working closely with AI and data teams to ensure that all AI initiatives comply with relevant laws and regulations. This collaboration ensures AI systems are not only effective but also ethical and legally sound.

The primary goal of a compliance-centric approach is to achieve compliance and risk management by thoroughly documenting and attesting to AI use. This involves maintaining detailed records of data sources, model development processes and decision-making criteria to ensure transparency and accountability.

Why AI governance is essential

AI presents enormous potential. But its transformative power also comes with significant responsibility. Harnessing the full potential of AI while safeguarding ethical standards and public trust, includes a focus on ensuring AI systems are:

- Reliable
- Traceable
- Compliant

At Collibra, we understand the complexities and challenges that come with deploying AI at scale. That's why we emphasize the importance of AI governance — a comprehensive approach to managing the AI use case lifecycle. By implementing robust governance practices, organizations can navigate the risks associated with AI, ensure data quality, foster transparency and comply with evolving regulations.

Reliability: Building trust

Delivering reliable AI that users can trust is essential. By maintaining high standards of data integrity and model accuracy, your enterprise can build trust in AI systems for both internal and external users. Collibra helps ensure the reliability of AI models and the data that feeds robust governance capabilities, enterprise data catalog and continuous data quality checks.

Traceability: Ensuring understanding

Transparency in AI is key to building trust and understanding. Collibra provides a purpose-built AI governance solution that connects data, policies and models, enabling users to critically examine and explore the lineage of all AI inputs and outputs. With Collibra, you can build externally referenceable AI model cards, ensuring AI processes are transparent and understandable.

By ensuring transparency through traceability, Collibra helps organizations build AI systems that are not only reliable but also understandable to all users. This transparency is fundamental to gaining the trust of customers, regulators and other stakeholders, and it is a key component of responsible AI governance.

Compliance: Empowering stakeholders

Navigating the complex landscape of AI and data regulations requires clear insight and control. Collibra empowers AI stakeholders, including Risk and Legal teams, to manage and mitigate risks, protect data and demonstrate compliance with all relevant laws and regulations. Our comprehensive view across your entire data and AI infrastructure simplifies the compliance process, making it easier to adhere to both current and emerging standards.

Significantly, Collibra facilitates collaboration across stakeholder teams, including data scientists, legal counsel and compliance officers. It is this collaborative approach that ensures all perspectives are considered when developing and implementing AI governance strategies. By fostering open communication and shared responsibilities, we help organizations build a culture of compliance that permeates every level of the organization.

AI and ethics: Good for society, good for business

As AI continues to transform industries, the importance of responsible governance cannot be overstated. AI and data have the power to reshape our world, from healthcare to government to social media platforms and much more.

However, with this power comes the responsibility to address ethical challenges such as bias, privacy and fairness. At Collibra, we believe that ethical AI governance is crucial for ensuring that AI innovations benefit society as a whole.

By implementing AI governance, organizations can:

- Mitigate risks: Identify and manage potential risks associated with AI, from data privacy concerns to regulatory compliance issues
- Enhance data quality: Ensure that the data used for training AI models is of high quality, reducing the risk of biased or inaccurate outputs
- Foster collaboration: Promote cross-functional collaboration among data scientists, legal teams, and business stakeholders to create AI models that are ethical and effective

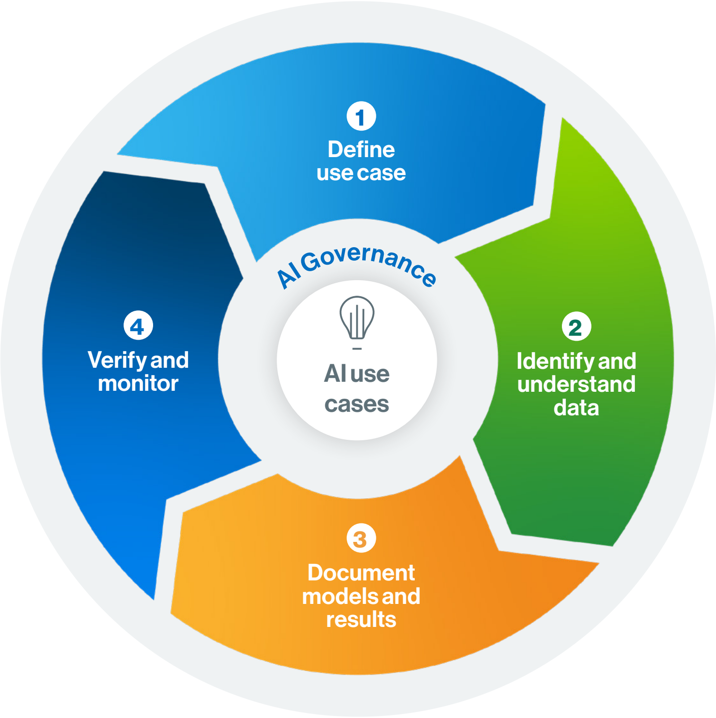

Our AI governance framework

To navigate the challenges to AI success, Collibra has developed a comprehensive solution based on the Collibra AI Governance Framework.

The Collibra AI Governance Framework provides a structured approach to managing AI initiatives, ensuring that AI is used responsibly, maximizing value while minimizing risks. The Framework includes the following four steps:

- Define the use case
- Identify and understand data
- Document models and results
- Verify and monitor

By following these steps, you can harness the full potential of AI, driving innovation and achieving significant competitive advantages.

Step 1: Define the use case

The first step in your AI journey is to clearly define your use case. Knowing the intended purpose of your AI model — and where it will be deployed — should always be your first step.

A well-defined use case serves several purposes. It clarifies why the AI model is necessary, outlines the specific problems it aims to solve and details the type of data it will utilize.

Collibra AI Governance: The solution for every AI use case

Collibra AI Governance offers a holistic end-to-end view of all your organization's AI projects, and is designed to help improve organizational access and adoption of AI systems. Collibra AI Governance is essential to promoting visibility, productivity, compliance and accountability for AI use cases.

What goes into defining a use case?

- Business context: Develop a well-documented use-case description that includes an analysis of the business value, the business policies the model may impact, and a list of business owners and their respective responsibilities
- Legal, ethics and compliance: Assess whether the model will handle sensitive or private information,
  such as personal identifiable information (PII). Understand and document any specific regulations that may impact your AI model, along with risk assessments to ensure compliance and ethical considerations are addressed
- Data usage: Clearly outline the data required for the model, including what data will be used as input, how the model will be trained and the nature of the output data
- By addressing these key areas, you'll ensure your AI roadmap is grounded in a thorough understanding of the broader context before you invest any resources into building an AI model.

Step 2: Identify and understand data

AI starts with trusted data. It's why the old adage — "Garbage in, garbage out" — is especially true when it comes to AI. It explains why once you've defined your AI use cases, you need to take a close look at your data.

But how? Most enterprises manage legions of databases in a complex data landscape. It's why the cornerstone of any successful AI initiative is a deep understanding of the data your model will leverage. It includes understanding the nature of your data, as well as ensuring your compliance with all relevant laws and regulations. And it's why delivering trusted data for AI models starts with implementing the right data governance strategy. Rigorous guardrails will ensure you can operationalize AI workflows and processes to deliver trusted data.

To operationalize successfully, an enterprise data catalog is key. It streamlines discovery and understanding of data across sources. Plus, Collibra includes a user-friendly data marketplace, in addition to a data catalog, that helps data scientists find and access data in a fraction of the time compared to traditional methods of consulting stakeholders.

Additionally, Collibra provides active data pipeline monitoring using advanced data quality and observability tools to quickly identify and resolve problems before they reach your AI models.

Step 3: Document AI models and results

With a well-defined use case and high-quality data to feed your model, your focus shifts to building the AI model. It's crucial to document every detail during this process, including model outputs and challenges faced.

This step is where data scientists will focus most of their time. They'll document, trace and track the model, associated data products and usage. Comprehensive documentation is vital for model analysis and reporting. Data lineage is particularly essential in this phase; it ensures that you have clarity on the origin of the data, any transformations to it, and how and where outputs are used. This is especially useful in highly regulated industries, like financial services, where regulators may demand to see how data is being used.

In this step, your primary goal is to get initial results. Once you land on a model that passes scrutiny, you're ready for the hard step of moving into production.

Step 4: Verify and monitor

It's important to remember that AI governance is not a one-time effort. Once your model is ready for production, it's vital you continually monitor results and revisit the legal and compliance requirements as new AI- and data-specific regulations are always coming into play.

The key aspects of this step include:

- Verifying model performance: Prior to full-scale deployment, it's critical to verify that the AI model acts as intended. Verification is a quality check, confirming that the model meets technical and business expectations
- Putting the model into production: Moving the AI model from a controlled testing environment into production is a big step. It involves integrating your model into your operational environment where it will start affecting real-world decisions. You'll also trace and document the flow of data through the AI system to understand how data is transformed and used in decision-making, which is crucial for troubleshooting and compliance
- Ongoing monitoring for data quality and compliance: Monitoring is vital to detect and address performance issues, data drifts or unexpected behavior. It involves tracking model output for accuracy, bias, and adherence to regulatory and ethical standards. You'll also vigilantly protect sensitive data, adhering to privacy regulations and ethical standards, especially as the model interacts with new datasets
- Retraining the model as needed: AI models are not set-and-forget tools. They require periodic retraining to incorporate new data, new regulations, and new technologies. Retraining is crucial to ensuring the model's accuracy and relevance

By following this step with an emphasis on data quality, data lineage and data privacy, you can ensure your AI models remain relevant, robust and compliant, as well as capable of adapting to new challenges and ensuring long-term effectiveness.

AI governance: Critical capabilities checklist

Researching AI governance solutions? A comprehensive AI governance platform that focuses on ensuring trusted data requires an array of advanced tools. Here are the key capabilities to look for in an AI governance solution:

- Data Governance - Automate data governance workflows to deliver trusted data to AI teams faster
- Data Catalog - Help AI teams discover and understand data they have for performant AI
- Data Quality & Observability - Ensure high-quality data in every model built
- Data Lineage - Map input, training and output data for analysis and regulatory reporting
- Data Privacy - Automate and operationalize privacy and address both global AI and data regulations
- Data Access Governance - Classify and protect data across all data sources

Conclusion

Collibra AI Governance: Your competitive advantage

Today, AI governance is not just a compliance checkbox—it's a competitive advantage and a strategic imperative. Collibra AI Governance is the leading governance framework for building trust, ensuring compliance, maintaining security and guaranteeing reliability for AI initiatives. By partnering with Collibra, your organization can confidently navigate the complexities of successfully deploying your AI projects while mitigating risks, building stakeholder trust and unlocking the full value of AI.

Take a tour of Collibra AI Governance
Schedule a demo