AI Governance 101: The First 10 Steps Your Business Should Take

Insights

10.31.24

Artificial Intelligence (AI) is reshaping how your business operates – but adopting AI technology comes with new responsibilities. The only way to ensure that the AI you use aligns with your business goals, complies with regulations, and operates ethically is through "AI governance." This 10-step guide provides you with straightforward steps for getting started, covering the basics and helping you avoid risks from day one. *And if you'd like to learn more, make sure to sign up for our complimentary Standing Up AI Enterprise Governance webinar on November 12.*

1. Understand What AI Governance Is

Let's start with the basics to make sure we're on the same page. AI governance is simply a framework of rules, practices, and policies that direct how your organization manages AI technologies. Governance is about building a process, following the process, and documenting the process. These steps will ensure that your AI technology aligns not only with your company values and customer expectations but also with the legal standards that are springing up as courts and government investigators adopt them. The five key principles to keep in mind when building this structure are:

- **Transparency**: Ensure that AI decisions are explainable and understandable.

- **Fairness**: Avoid unintentional biases that can lead to discriminatory outcomes.

- **Accountability**: Designate responsibility for AI outcomes to humans.

- **Privacy and Security**: Protect user data and adhere to federal and state data protection laws.

- **Document**: Document every step in the AI governance process. If it's not in writing, it didn't happen.

What are the Principles of AI Governance?

TRANSPARENCY
Ensure that AI decisions are explainable and understandable.

FAIRNESS
Avoid unintentional biases that can lead to discriminatory outcomes.

ACCOUNTABILITY
Designate responsibility for AI outcomes to humans.

PRIVACY AND SECURITY
Protect user data and adhere to federal and state data protection laws.

DOCUMENT
Document every step in the AI governance process. If it's not in writing, it didn't happen.

Fisher Phillips

2. Form an AI Governance Committee

Your first real step is assembling the right team to oversee your efforts. Create a cross-functional AI Governance Committee that includes representatives from legal, IT, human resources, compliance, and management. This team will oversee your organization's AI implementation, monitoring, and auditing.

3. Define and Document AI Use Cases

Before deploying any new AI system or software, clearly define and document its intended use. For each AI initiative, clarify the following:
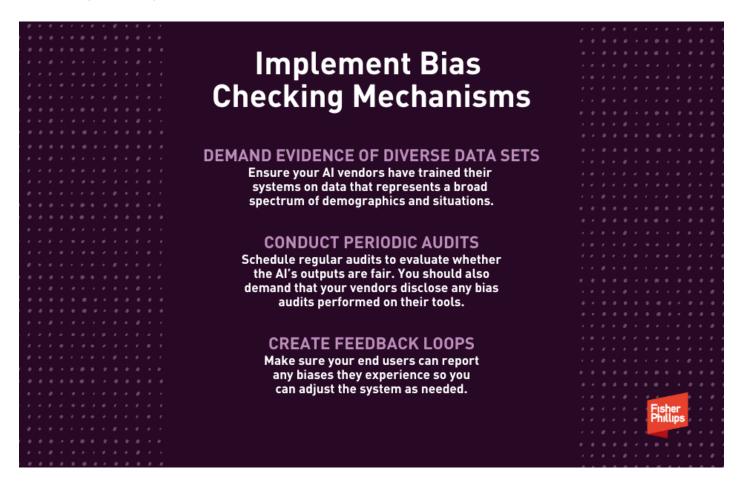
- **Purpose**: What business need does this AI solution serve?

- **Data**: What data will it use, and how will data be collected and protected?

- **Ethical and Legal Boundaries**: What ethical or legal issues do we need to consider?

Answering these questions – and clearly documenting them – will create a trail that makes your AI actions traceable and ensures accountability. Should questions arise, documented use cases can clarify AI's role in your operations.

4. Implement Bias-Checking Mechanisms

One the biggest AI dangers that is on the minds of government investigators and plaintiffs' attorneys is the concern about unintentional discriminatory bias. You probably know that AI systems can unintentionally reinforce biases that were present in the data they're trained on, but you might not know that most courts and agencies won't let you easily escape liability by pointing your finger and blaming your AI vendor for causing such discrimination. That means you need to regularly review and test your AI models for bias to prevent unintended outcomes. Some practical steps you can take:

- **Demand Evidence of Diverse Data Sets**: Ensure your AI vendors have trained their systems on data that represents a broad spectrum of demographics and situations.

- **Conduct Periodic Audits**: Schedule regular audits to evaluate whether the AI's outputs are fair. You should also demand that your vendors disclose any bias audits performed on their tools.

- **Create Feedback Loops**: Make sure your end users can report any biases they experience so you can adjust the system as needed.

Implement Bias Checking Mechanisms

DEMAND EVIDENCE OF DIVERSE DATA SETS
Ensure your AI vendors have trained their systems on data that represents a broad spectrum of demographics and situations.

CONDUCT PERIODIC AUDITS
Schedule regular audits to evaluate whether the AI's outputs are fair. You should also demand that your vendors disclose any bias audits performed on their tools.

CREATE FEEDBACK LOOPS
Make sure your end users can report any biases they experience so you can adjust the system as needed.

Fisher Phillips

5. Establish Clear Accountability Pathways

Next, make sure to assign responsibility for AI outcomes to specific roles within your organization. This not only is a good best practice but also can protect your organization from risks. Define who oversees AI systems, who is accountable if the AI doesn't perform as intended, and who manages data security. Some possible examples could include:

- **Data stewards** can be responsible for data quality and protection.
- **Algorithm auditors** can regularly review algorithms for performance and ethical alignment.
- **Compliance officers** can ensure your AI use complies with regulations or government guidance.

6. Use Real-World Scenarios to Create Practical Guardrails

Implementing AI without considering possible missteps is a mistake. Run through some possible mistakes or errors that your AI platforms could cause in worst-case scenarios and gameplan some guardrails you can put in place to prevent them. (This is a good exercise for ChatGPT or Claude, by the way.) Some examples:

- *Scenario 1: Discriminatory Hiring Algorithms*
  What if your AI system that screens resumes has an unintentional bias and favors certain demographics? Solution: Make sure your AI is trained on diverse data sets, but also conduct regular bias checks to monitor it.

Scenario 1: Discriminatory Hiring Algorithms

What if your AI system that screens resumes has an unintentional bias and favors certain demographics?

Solution: Make sure your AI is trained on diverse data sets, but also conduct regular bias checks to monitor it.

- *Scenario 2: Problematic Chatbots*
  What if your AI-powered chatbot (a probabilistic tool that replaced your deterministic chatbot) gives inconsistent responses and frustrates customers? Solution: Regularly monitor interactions for accuracy and fairness – and make sure there's a way for users to report issues through clear feedback loops.

Scenario 2: Problematic Chatbots

What if your AI-powered chatbot gives inconsistent responses and frustrates customers?

Solution: Regularly monitor interactions for accuracy and fairness – and make sure there's a way for users to report issues through clear channels.

- *Scenario 3: Overeager Predictive Maintenance Alarms*
  What if your AI model that predicts equipment breakdowns sends frequent false alarms and disrupts your operations? Solution: Implement regular feedback loops with operators to fine-tune the model to ensure efficiency and accuracy.

Scenario 3: Overeager Predictive Maintenance Alarms

What if your AI model that predicts equipment breakdowns sends frequent false alarms and disrupts your operations?

Solution: Implement regular feedback loops with operators to fine-tune the model to ensure efficiency and accuracy.

- *Scenario 4: Targeted Marketing Gone Awry*
  What if your AI recommendation engine sends inappropriate product recommendations (like advertising luxury goods to low-income users or cleaning products only to women)? Solution: Avoid assumptions based solely on limited demographic data and refine recommendations based on customer feedback.

Scenario 4: Targeted Marketing Gone Awry

What if your AI recommendation engine sends inappropriate product recommendations (like advertising luxury goods to low-income users or cleaning products only to women)?

Solution: Avoid assumptions based solely on limited demographic data and refine recommendations based on customer feedback.

- *Scenario 5: Financial Risk Management is Too Risky*
  What if your AI-driven financial model makes risky investment recommendations that inadvertently breach compliance guidelines? Solution: Use human oversight to periodically audit and approve AI-generated recommendations to ensure regulatory alignment.

Scenario 5: Financial Risk Management is Too Risky

What if your AI-driven financial model makes risky investment recommendations that inadvertently breach compliance guidelines?

Solution: Use human oversight to periodically audit and approve AI-generated recommendations to ensure regulatory alignment.

7. Conduct Regular Audits and Training

AI is constantly evolving, and so are the rules that govern it. Schedule regular audits to evaluate your AI's alignment with the latest governance principles and best practices. Equally important, train your employees to understand AI basics, governance policies, and ethical considerations on at least an annual basis.

8. Document and Track Decisions

For effective AI governance, documentation is essential. Keep records of:

- **Model Versions**: Track which versions are in use and any updates applied.

- **Policy Changes**: Log updates to AI governance policies, documenting the reasons and any associated risks.

- **Decision Rationale**: When AI makes significant decisions, log the rationale to ensure traceability and accountability – and note the extent to which human judgment came into play before the decision was implemented.

9. Stay Informed on AI Governance Trends

AI governance is an emerging field, with new regulations and best practices developing quickly. Make a habit of staying informed on industry news, emerging standards, and evolving regulations. Some effective ways to stay updated include:

- **FP Insights**: Subscribe to FP Insights and make sure to click the "AI, Data, and Analytics" box to receive AI-related updates.

- **FP Events**: Sign up for our complimentary *Standing Up AI Enterprise Governance* webinar on November 12 and stay tuned for further events.

- **Professional Networks**: Join AI-focused groups or industry associations to exchange knowledge.

- **Training**: Encourage team members to attend webinars and workshops.

- **Partnerships**: Work with AI governance specialists who can guide you in compliance and risk management.

10. Consider a Trusted Partner

Navigating AI governance can be challenging, especially for businesses new to AI. Having a trusted legal or consulting partner can provide valuable guidance, ensuring compliance and aligning AI practices with business goals. From setting up policies to managing compliance, a partner experienced in AI governance can help prevent costly missteps. Your FP attorney can connect you with a member of the FP Artificial Intelligence team to assist.

Conclusion: Start Small, Scale Smart

Establishing effective AI governance takes planning, but these foundational steps can set you on the right path. We'll continue to monitor developments in this ever-changing area and provide the most up-to-date information directly to your inbox, so make sure you are subscribed to Fisher Phillips' Insight System. If you have questions, contact your Fisher Phillips attorney, the authors of this Insight, or any attorney in our AI, Data, and Analytics Practice Group.

Related People

Richard R. Meneghello
Chief Content Officer
503.205.8044
Email

David J. Walton, CIPP/US
Partner
610.230.6105
Email

Service Focus

AI, Data, and Analytics

Privacy and Cyber