How to govern your AI strategy in a highly regulated world

Keep your organization and the people who interact with it safe and responsible in the face of growing AI regulations. And gain the freedom and confidence to innovate with agility.

sas

Contents

1.0 The growing importance of AI governance

AI has become fundamental to transforming processes in most organizations.

From accelerating innovation by integrating AI models into decision workflows to recognizing customer patterns or optimizing marketing strategies. It's also being used to enhance threat detection by spotting unusual patterns or behaviors and automating repetitive, manual tasks to free up valuable resources. These are just some examples of how AI is redefining work in material and scalable ways.

In society, we're seeing AI help doctors diagnose medical conditions earlier and enable innovation in medicine. The technology is even being used to support conservation efforts of threatened species and habitats. For example, SAS is working with the UNC Center for Galapagos Studies to help protect endangered sea turtles, using crowd-driven AI to help train a SAS computer model to identify sea turtles using facial recognition.[1]

AI is transforming the workplace and society at large at a pace that stands to catch us off guard if we're not fully prepared.

Generative AI is a subset of AI and is quickly catching on in organizations around the world.

While traditional AI and machine learning systems recognize patterns in data to make predictions, GenAI goes beyond by generating new data as its primary output. It can generate music, art or images from text-based descriptions or even develop a business strategy with conversational, back-and-forth prompts.

Coleman Parkes conducted a survey on behalf of SAS in April 2024, targeting 1,600 decision makers in GenAI strategy or data analytics in organizations across key sectors globally.

The research shows that those organizations embracing GenAI have seen many benefits, with 89% citing improved employee retention and satisfaction and 82% seeing operational cost savings.[2] Early successes are likely to drive exponential growth in the use cases of GenAI.

But there is still work to be done, with only 5% of organizations surveyed confident that they have a comprehensive governance framework for AI, demonstrating a huge gap in readiness and exposing those unprepared organizations to indeterminable risk.[2]

of organizations feel confident they have a comprehensive governance framework for AI.

Governance is essential to the reliable, transparent and ethical performance of AI. But the clear lack of oversight of the technology is putting many organizations at risk of noncompliance with incoming regulations around AI, such as the EU AI Act.

[1] https://www.sas.com/en_us/news/press-releases/2023/may/unc-center-for-galapagos-studies.html
[2] Generative AI Global Research Report: Strategies for a Competitive Advantage: https://www.sas.com/content/dam/SAS/documents/marketing-whitepapers-ebooks/ebooks/en/generative-ai-global-research-report-113914.pdf

2.0 What does AI governance look like?

AI governance is a culture guided by values and principles designed to ensure the ethical, safe and effective use of AI. It's based on both an organization's internal motivations and its external pressures. It is an all-encompassing strategy that establishes oversight, ensures compliance, and develops consistent operations and infrastructure within an organization.

With AI governance in place, as an organization's adoption of AI grows and evolves, it is protected with oversight, enhanced compliance, and consistent operations and infrastructure – fostering a culture of trustworthy AI.

Values indicate what the stakeholders care about and want to protect and promote.

Principles are the norms or guardrails that describe what is ethically and legally required to protect and promote the **values.**

Governance assesses if the **principles** are satisfied in a specific case.

Ethics is reflected in the labor performed to uphold **values**, adhere to **principles** and support **governance**.

3.0 Why it's more important now than ever

It's no longer a case of "if" an organization adopts AI but "when." Those who don't stand little chance of keeping up with their industry or the competitors that are using the technology; those deploying AI solutions can propel innovation, insights and beyond: they can see more, and do more, with fewer resources.

But organizations that are powered by AI move fast, and computer scientists have known for decades that AI systems drift over time and need to be monitored. So, it's crucial that guardrails are put in place to protect organizations and citizens alike; without these checks and balances, the risk of unethical AI is very real.

With thorough AI governance in place, organizations get a faster onramp to new technologies, improved agility and the ability to innovate. GenAI can help drive product innovation and identify and create new business opportunities – leading to new revenue streams and potential growth. With the right governance, the adoption of new technologies is smoother, and there's less likelihood of problems caused by those new technologies downstream.

So, before embarking too far down the path to AI, it's crucial to consider the implementation of a clear AI governance strategy from the start. This way, you can avoid the wrong turns, blind bends and potholes that your hastier peers might encounter.

The 5 key drivers for implementing AI governance today:

1. Regulatory compliance

With governments around the world concerned about the implications of AI in their respective regions, many are moving to put regulations in place. Those organizations that establish strong AI governance – even going above and beyond the requirements set out – are going to be in a strong position to continue their pace of progress as new regulations and guidance start to take a stronger hold, as they will already have important infrastructure in place to eventually meet requirements.

2. Productivity

AI talent can be the costliest for organizations. With clear guidelines grounded in values and guided by principles, you can enable distributed decision making – which, in turn, supports increased productivity, as employees can make decisions more confidently and sign-off becomes faster. Practicing governance also allows for the uptake of new technologies with greater speed and agility.

3. Win and keep talent

The majority of AI talent want to do good with their skill set and be sure they're improving the lives of others. Organizations that can signal a highly functioning AI governance program are more likely to attract great talent because they'll feel safe that their work will be used responsibly and to help people.

4. Trust

By making clear and measurable commitments to which your organization can be held accountable, your brand will see a boost in trust among stakeholders. And by failing to meet the guidelines, you'll be at risk of huge monetary and reputational loss.

5. Increase brand value

Customers care about the impact an organization has on the environment and society. And as more information is made accessible through social media and other digital channels, there's no hiding messy ethics. With strong AI governance safeguards in place, you can avoid, or more quickly recover from, unintended societal and environmental harm.

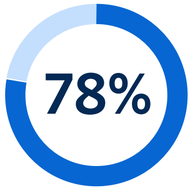

believe that organizations that utilize GenAI in their business or activities are extremely or very responsible for ensuring that GenAI is developed and used ethically.[3]

[3] https://kpmg.com/kpmg-us/content/dam/kpmg/corporate-communications/pdf/2024/kpmg-generative-ai-consumer-trust-survey.pdf

4.0 The risks of poor AI governance

The impact of AI technology can reach far and wide. Like a physical supply chain, it is crucial to manage and maintain an end-to-end view of the entire AI life cycle, from inputs to outputs.

Transparency and explainability should extend as much as possible to those subjected to AI systems – whether they are aware of it or not – and provide easy, intuitive ways for them to seek recourse. It is also important that citizens have visibility, control and the ability to give approval when appropriate.

AI can pose a significant risk if it is developed, deployed and managed without intentionality and discipline. For example, AI systems trained on biased data may learn and replicate historical patterns of discrimination against women, people of color or other vulnerable groups. If these biases are not identified and mitigated before deployment, AI systems could perpetuate social biases and lead to unintended large-scale consequences.

Then there are, of course, the financial implications and reputational risks of AI that are not correctly governed. The EU AI Act has fines that can reach 7% of worldwide revenue.

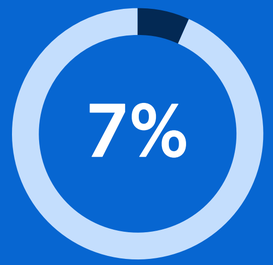

of worldwide revenue
... the potential fine for breaching the EU's AI Act

5.0 Growing regulatory pressure

We've touched upon the risk of monetary fines – not to mention the brand damage – of falling foul of new AI regulations that are coming into force. It's a key driver behind AI governance and the growing role it plays in protecting organizations by placing safeguards around their use of AI.

Let's take the EU AI Act as an example of a regulation that's been created to avoid the potentially harmful consequences of the unmanaged use of AI and ensure its safe and ethical use. By adhering to the guidance and rules set out by the EU AI Act, organizations that do business in or with the EU can build trust in their use of AI through accountability.

EU AI Act – what is it?

The goal of the EU AI Act is to lay down a harmonized legal framework "for the development, the placing on the market, the putting into service and the use of artificial intelligence systems" in the EU. It's extensive, spanning 180 recitals and 113 Articles, the new law takes a risk-based approach to regulating the entire life cycle of different types of AI systems.

What's at stake?

Noncompliance with the EU AI Act will be met with a maximum financial penalty of up to EUR 35 million or 7% of worldwide annual turnover, whichever is higher.

Does it apply to non-EU organizations?

It does. The Act applies to any non-EU company that offers products or services within the EU market. So, any company, regardless of its geographic location, must comply with the EU AI Act if it targets EU users or affects EU citizens.

A global approach to governing AI: The UN's AI Advisory Council

In 2023, the UN formed an AI advisory body to undertake analysis and advance recommendations for the international governance of AI. In September 2024, the Council published its **final report**, of which the focus was the clear need for global governance to address the challenges and risks of AI while also ensuring that the full potential is harnessed in ways that no one is left behind.[4]

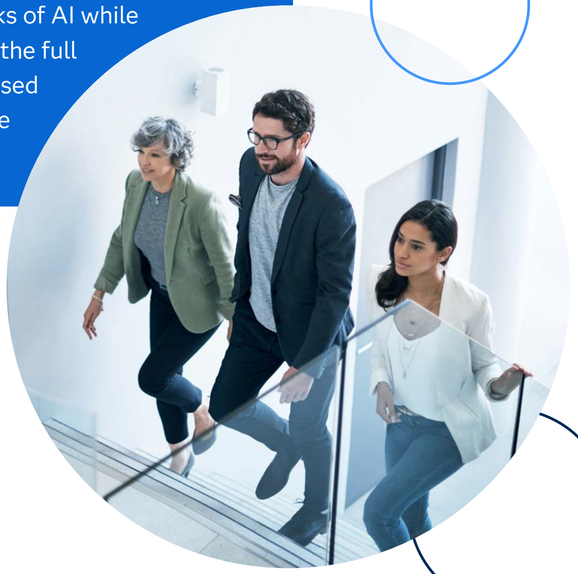

[4] UN Governing AI For Humanity: **https://www.un.org/sites/un2.un.org/files/governing_ai_for_humanity_final_report_en.pdf**

6.0 How do you measure readiness?

With solid AI governance processes and frameworks in place and the technology to support them, organizations won't be phased by typical ethical questions about AI. Employees are confident in their ability to raise a flag if they notice anomalies, deviation from norms or a risk that needs attention.

> **"One thing is true globally, though, and that's the fact that all governance activity associated with AI only moves as fast as the speed of culture."**

Reggie Townsend, Vice President of the Data Ethics Practice, SAS.

What's your risk appetite?

Some organizations find it useful to implement a risk-based approach to AI development and usage. Understanding an organization's risk appetite, tolerance and threshold may assist in establishing corporate strategy, policy and oversight of the design, development, deployment and use of AI systems. And regulators are embracing risk-based approaches.

Risk models

There are multiple approaches to formulating risk models. One is rating the likelihood of an event happening and the severity of impact if it does. With this approach, organizations would focus on those scenarios with the most severe impact that have the highest potential of happening.

Another model is to look at risk the way trust and safety groups often do, asking, **"Is this an attack on individuals, an attack on an organization or an attack on our way of life?"**

Criteria for assessing readiness

There are four key pathways to maturity for AI governance: **oversight, compliance, operations and culture.**

Oversight and **compliance** are where much of the governance and most of the legal work happens. **Operations** is where most technical interventions and tools appear. **Culture** is the piece that makes everything possible, including maintaining governance over time.

In summary: Best practice guidance for implementing a successful AI governance strategy

Operationalize – Train the workforce, gain regulatory insight and create a decision framework.

Document – Establish principles, write policies and communicate plans.

Engage – Involve stakeholders early, take part in compliance activities and work with industry consortia.

7.0 The SAS approach to AI governance

Driving awareness of the need for AI that is ethical, fair and sustainable is central to everything we do at SAS.

We always ask not just “Could we?” but also “Should we?” when considering innovation. Our principles are reflected in our people, our processes and our products. These principles help everyone at SAS navigate, manage and negotiate ethical tensions in the most productive and least harmful ways feasible.

SAS’s six core principles for responsible innovation:

Human Centricity Promote human well-being, human agency and equity.	**Inclusivity** Ensure accessibility and include diverse perspectives and experiences.	**Accountability** Proactively identify and mitigate adverse impacts.
Transparency Openly communicate intended use, risks and how decisions are made.	**Robustness** Operate reliably and safely, while managing potential risks throughout a life cycle.	**Privacy & Security** Protect the use and application of an individual’s data.

From the SAS Data Ethics Practice was born a collaborative methodology that focuses on four pillars: Oversight, Operations, Compliance and Culture. Together, these help us to anticipate, mitigate and avoid unintentional harm, particularly for the most vulnerable.

Oversight

Ensures that internal and external AI-enabled technologies and processes adhere to data ethics principles. Advises leadership on sales, consulting, product development and procurement opportunities involving AI.

Operations

Absorbs and interprets expectations and demand, then develops AI technologies that align with market viability, SAS portfolio synergy, regulatory compliance and data ethics principles.

Compliance

Monitors, audits and seeks compliance with data ethics principles. Provides organizational checks and balances to identify mitigation needs before AI-related technologies and services are made public.

Culture

Cultivates an ecosystem of creators, contributors and consumers of AI for knowledge sharing, collaboration and the normalization of behaviors and practices aligned with data ethics principles.

AI governance maturity model

When it's embedded into your company culture, you're well on the way to effective governance.

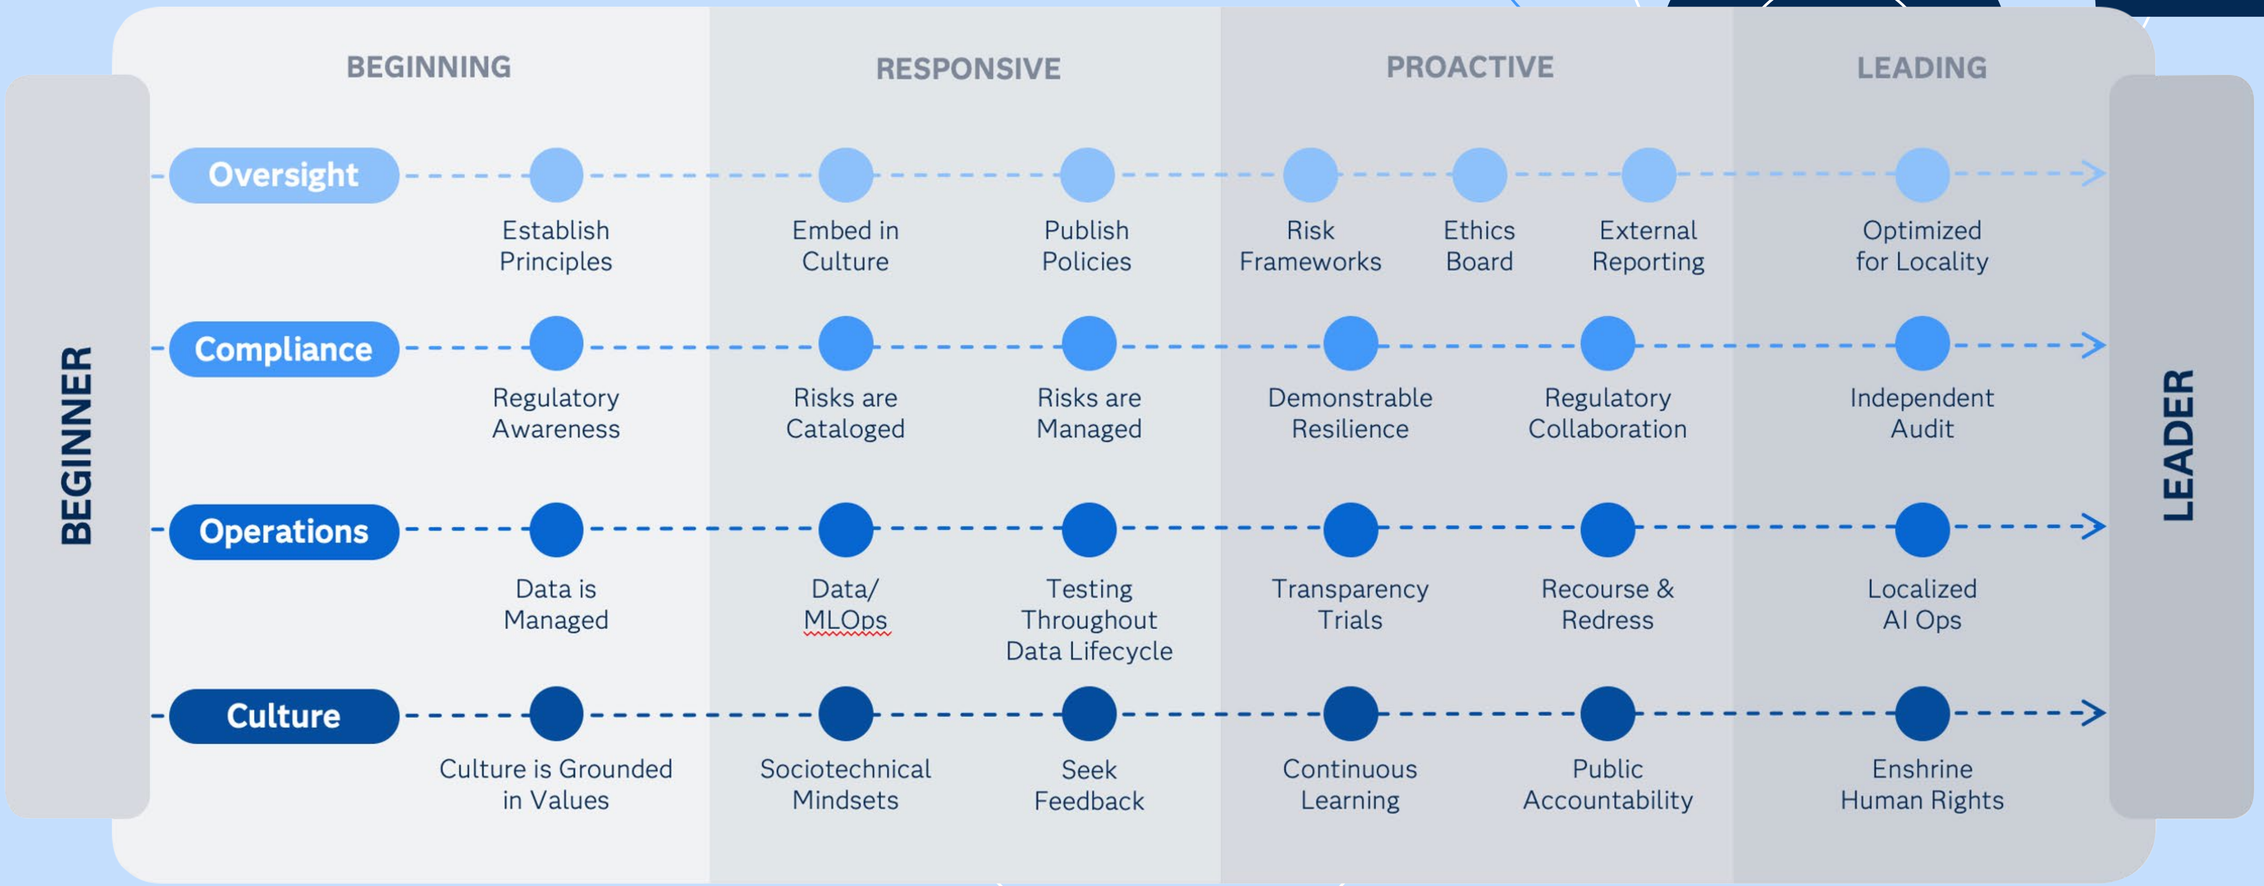

A practical approach: Implementing AI governance with SAS

When you work with SAS, you get a proven approach that combines AI governance with principle-based technology. So, you can ethically and responsibly unlock the potential of AI in your organization with a trusted, end-to-end approach across your AI life cycle. This allows your organization to:

Use AI ethically with a summary of a model's training data, intended use and performance – demonstrating whether models are aligning with ethical AI goals. Transparency and explainability enable a proactive approach to understanding and mitigating risk.

Consider challenging AI regulations, risk and governance with support from AI experts on oversight, compliance, operations and culture. With practices relevant to any new regulations that may come tomorrow and prioritizing privacy and security with information and model governance.

Build consistency and trust with explainable AI and data lineage that facilitates human-in-the-loop transparency. And curate and govern LLM prompts using natural language processing to improve accuracy and repeatability.

Embed governance throughout the data and AI lifecycle with built-in machine learning and AI capabilities to scan data and make intelligent suggestions, enabling automated workflows or the use of synthetic data. Reduce bias with fairness assessments and bias mitigation, while performance monitoring, model drift and interpretability provide insight into the results of a model.

Trustworthy AI across the data and AI life cycle

A drive to create. A duty to care.

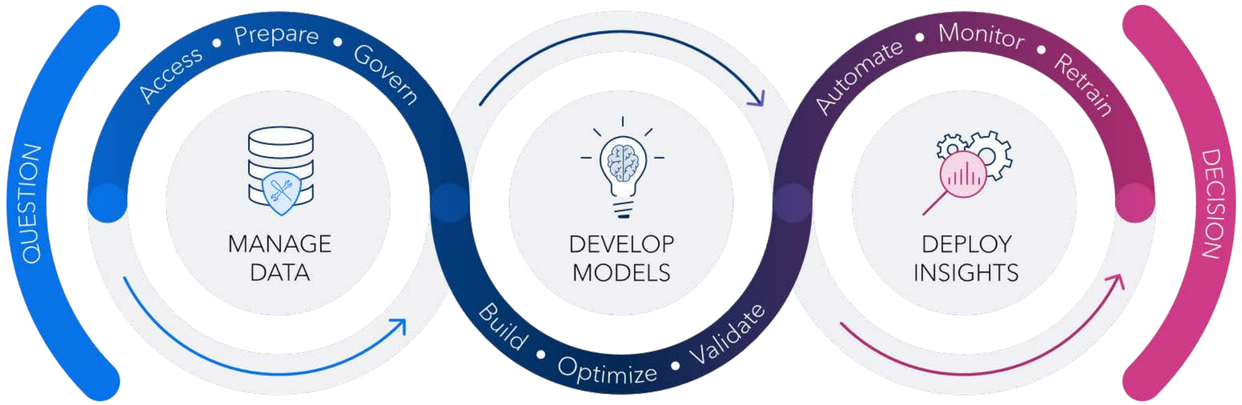

Data flows into and through an organization in numerous ways. If you're not vigilant, biases can build up along the way, leading to decisions that do more harm than good.

Tackle any data and AI challenge with robust data governance, transparent AI and secure ModelOps processes on the SAS® Viya® platform.

SAS Viya employs a trustworth AI approach by injecting fairness and oversight at every step of the data and AI journey so you can generate outcomes that are ethical, equitable and sustainable.

Understanding where your organization is on its journey to AI governance.

Stage 1: Unaware	**Stage 2:** Fragmented	**Stage 3:** Top-down	**Stage 4:** Pervasive	**Stage 5:** Steward
AI governance is something other organizations do and is not part of the organization's strategy or risk management process. The organization has not provisioned governance measures for AI, nor does it have plans to do so.	Leadership sees AI governance as a potential option for the future. Most executives don't see AI as a risk. There could be isolated pilots for governance. Leadership does not have AI governance in its current plans.	The CEO has set the tone for AI governance, initiated a top-down initiative, and expects to capture a competitive advantage from these efforts. Working on these initiatives is sought after by the best talent.	C-suite owns AI governance and policies, training, frameworks, and reporting are commonplace. AI governance is routine, part of the culture and predictable.	Connected to an ecosystem of peers, sharing information is commonplace. Employees demonstrate competence, share thought leadership and contribute to standards. Some provide governance services.

8.0 A closing note

Using this guide as a starting point for developing a rigorous AI governance policy, organizations can be ready for whatever is around the corner in the complex terrain that is AI and GenAI.

A solid governance strategy will not only bring peace of mind to the organization and the freedom to innovate with protective guardrails in place but will also build a foundation for trust among customers and stakeholders, taking into account regulatory compliance organizational needs along the way. With all this taken care of, organizations are free to use AI to innovate for good, push boundaries and surpass the competition.

With SAS, you get a proven approach and technology to make your AI governance strategy more than words on a page. You'll get the tools to support your strategy across the entire AI and analytics life cycle. And it's all reinforced with learnings and insights from our Data Ethics Practice to ensure you're always using your data in the right way to stay compliant as AI takes the lead in transforming processes and data use.

Govern your AI strategy with confidence with SAS.

Learn more