




Creating an AI Governance Program

Abstract

The pressure for companies to deploy generative AI-assisted applications to gain a competitive advantage is steadily increasing. Yet the productivity and competitive advantages AI offers are countered by emerging compliance challenges and risks. Developing a comprehensive AI Governance program can enable faster AI deployment while also driving compliance, identifying, and managing risks and ensuring ethical and safe use.

Contoural Independence

Contoural is an independent provider of Information Governance services. We do not sell any products, provide any reactive eDiscovery services, or store any documents. We do not take any referral fees, including but not limited to products such as Microsoft 365. We generate our revenue from strategic consulting services, enabling us to develop and recommend strategies driven not by a product or eDiscovery agenda, but rather an approach based exclusively on the needs of our clients.

contoural 

Introduction

Pressure for companies to use generative Artificial Intelligence (AI)-assisted applications to gain a competitive advantage (or at least not fall behind versus competitors) is steadily rising, and in 2024, CEOs will push their IT, Legal, Compliance, and Privacy Teams to deploy AI applications now, not later. While AI promises tremendous innovation and productivity gains, the emerging compliance challenges and risks can feel overwhelming: Seemingly every week, new AI regulations are announced, the copyright and IP issues are only beginning to be addressed by the courts, and companies need to ensure they are using new AI technology both ethically and correctly. This business pressure to deploy AI countered by compliance, risks and accuracy issues is creating a tug-of-war within organizations.




Figure 1. The competitive advantage companies can achieve by deploying AI is countered by new and evolving compliance requirements, legal risks and ethics and safety issues, creating a tug-of-war on using this new technology.

There is a middle route that avoids deploying risky, non-compliant solutions or sitting on the sidelines as their competitors deploy AI. Smart companies today are developing AI Governance programs that drive compliance, identify, and minimize risks, and ensure the ethical, correct, and safe use of these technologies.

Benefits and Competitive Advantage Driving Explosion of AI

The advent of generative artificial intelligence (AI) offers the promise of tremendous leaps in productivity, new revenue, cost savings, and increased innovation. After decades of technological stagnation with respect to AI, generative AI has elevated it from the fringes to the mainstream. Contrary to doom forecasts from digerati of AI suddenly replacing (and taking over) segments of society, AI deployment will initially be piecemeal, with AI-assisted applications launched in legal, finance, marketing, product design, engineering, and eventually in nearly every single area of the organization. (Full disclosure: Contoural is launching an AI-based records management initiative.) Without overstating, AI has the potential to be transformative.

10X AI-Assisted Engineer

AI holds the promise of making individual workers ten times more productive, including workers in

- Software development
- Legal • Finance • Media
- Customer service

Competitive Advantage

THE WALL STREET JOURNAL.

Rush to Use Generative AI Pushes Companies to Get Data in Order

Data management is under the spotlight again as companies seek to out-innovate competitors with large language models

Stock Price Bump

"Companies that mentioned AI saw an average stock price increase of 4.6%, almost double that of companies that didn't."

— Research from
Storible and WaltStreet Zen

Figure 2. AI offers tremendous productivity increasing, driving competitive advantage.

Driving this rapid adoption is the promise of a 10X productivity increase, (called the "10X engineer"), which states that an employee leveraging an AI-assisted application can become ten times more productive. Software engineers use AI to develop code. Finance employees can use AI to automate many finance tasks and controls. In HR AI can dramatically reduce the time it takes to create job descriptions or update policies. Much of records management can be automated. The potential productivity impact of AI is real.

Longer-term AI has the potential to offer real competitive advantages, potentially reshaping how companies operate. Much as the advent of the internet in the late nineties changed many businesses operate, generative AI promises to be equally impactful.

Investors understand the tremendous competitive advantages AI offers and are rewarding companies that are or have announced they are using AI. According to researchers Storible and WaltStreetZen companies that mentioned AI saw an average stock price increase of 4.6%, almost double that of companies that did not. This in turn has created a concern that companies may be making false claims about their AI use. Recently Gary Gensler, Chair of the Securities and Exchange Commission, warned companies against "AI washing," or making false claims about their use of AI to drive up their stock price. Nevertheless, many CEOs in 2024 are eager to demonstrate that their companies are using AI.

Risks and Challenges of AI



**Emerging global
AI compliance and
privacy rules**



**Copyright and
intellectual
property/legal risks**



**Disclosure
of corporate
confidential
information and
trade secrets**



**Ethical use
of AI**



**AI accuracy,
correctness
and safety**

Figure 3. AI faces significant regulatory, legal, and other risks and challenges.

AI's tremendous benefits are being met with an almost equal concern on its risks. Regulators have been rushing to enact laws restricting how it can be used. Furthermore, the courts are only beginning to evaluate AI's copyright and intellectual property impacts. Finally, AI also raises significant questions on its ethical use as well as correctness and safety.


How to Think of Generative AI



AI Intern



**Human manager reviewing
work of a single AI intern**



**Human manager reviewing
work of multiple AI interns**

Figure 4. A knowledgeable, hard-working, but naive intern serves as a good analogy for AI.

In examining the risks and challenges of AI, it is important to understand its true capabilities. One analogy for generative AI's capabilities and limitations is as an intern. Imagine hiring a bright, hardworking, knowledgeable but sometimes naïve intern in the legal department. A risk with our AI intern is that they always want to please, and sometimes will fabricate information. With limited experience, this intern would not be given large or complex tasks. Rather, she or he would be given finite, specific assignments. Because the intern is new and inexperienced all their work product would need to be reviewed by their (human) manager. The manager would provide feedback, which the intern would use to continually improve their work. Once one intern has been successful, many more interns can be brought on board, scaling up the work on the given tasks, or having them tackle adjacent areas. Interns could even be hired to check the work of other interns. While much of the work would still need to be reviewed, a single human assisted by many interns would be much more productive, and even — with the right processes and quality feedback — higher quality work product.

Emerging AI Compliance

We expect to see many countries and states developing new rules limiting AI this year, creating a rushed and messy regulatory environment.

JURISDICTION	REGULATION	SCOPE
European Union	A.I. Act, GDPR Article 22	Comprehensive set of rules regulating the use of AI disclosure, consumer profiling, unfair discrimination in financial services and hiring, facial recognition, and registries. Expected to go into force in 2024 and apply in 2026. Additionally, GDPR prohibits automated decision making with personal information unless specific conditions are met.
US Federal Government	Executive Order (E.O.) 14110 on safe, secure, and trustworthy development and use of artificial intelligence	In addition to promoting safe use of AI, requires that AI developers share results of safety tests with regulators, provides guidelines for the federal government's own use of AI, and prohibits AI-driven discrimination.
US States	17 states	Since 2019, 17 states have enacted 29 bills focused on regulating the design, development and use of artificial intelligence. These bills primarily address two regulatory concerns: data privacy and accountability.
China	Various	China has released individual pieces of legislation whenever a new AI product becomes prominent. That's why China has one set of rules for algorithmic recommendation services (TikTok-like apps and search engines), another for deepfakes, and yet another for generative AI.
Rest of World	Globally bodies such as the UN, OECD, G20, and regional alliances have started to create working groups, advisory boards, principles, standards, and statements about AI. Likely to see growing differences between how democratic and authoritarian countries foster	

Table 1. Emerging AI regulations. This year will witness adoption of many new requirements.

Generative AI's explosive adoption has been met with a quick response from regulators. Every week governments across the world are proposing restrictions on how and where this new technology can be used. Wanting to become the global standard, European regulators announced restrictions on how AI can use information about individuals, as well as overall safeguards, especially around the use of personal information. In the U.S., states are limiting how companies can use AI to make financial decisions such as loan approvals. (Note that at least one data protection authority in the EU has also created such limits.) The Biden administration created a new standard for safety and security to protect privacy and civil rights. It can be argued that regulators are competing with each other and rushing to develop regulations, in hopes of setting a global regulatory standard. These new regulations are just the beginning, as we expect to see many countries and states developing new rules limiting AI this year, creating a rushed and messy regulatory environment.

FRAMEWORK	STATUS
AI Risk Management Framework	Currently a voluntary framework to incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems.
ISO/IEC 42001:2023 Information technology Artificial intelligence Management system	A comprehensive framework to ensure responsible AI development, deployment, and use, addressing ethical implications, data quality, and risk management.

Table 2. Emerging AI Frameworks

Copyright and IP Legal Uncertainty

In addition to new AI regulations, there are significant copyright and intellectual property concerns around AI. There are concerns that some “closed” AI Large Language Models from commercial vendors have been trained with copyrighted data. Furthermore, as these are closed systems, it is not possible to inspect what training data was used. Other generative AI vendors such as Adobe have gone out of their way to ensure that their products are based exclusively on fully licensed training data, even going as far as offering indemnification for copyright infringement claims for the users of their products. Ultimately, companies will have to determine their own level of risk tolerance.

In January, the New York Times sued both OpenAI and Microsoft claiming that their generative AI products were based on and violated the Times’ copyrighted information. The courts are just beginning to address some of these challenges, and we expect it may take years for instructive case law to provide any guidance. Companies deploying AI-assisted applications also need to ensure that these systems do not leverage or potentially expose proprietary, corporate confidential information or trade secrets.

In addition to the risk of misusing others IP, AI users also need to be careful about compromising their own IP or other sensitive data. For example, the Economist Magazine reported last year that Samsung employees unintentionally leaked proprietary source code via ChatGPT. An unprotected disclosure of a trade secret to a third party – through an AI assisted application – vitiates the status of the information as a trade secret.

Ethical Use of AI

AI systems are susceptible to the biases from the “training data” used to build the system’s intelligence, risking that they may lead to unethical actions. For example, if an HR application looking for candidates “teaches” an AI system to look for job candidates based on historical hiring profiles that do not reflect a company’s diversity goals, it may have an unintended bias. In this example, if the system is fed predominantly white males as examples to be used as the basis of “ideal” employees, the AI system may inadvertently only recommend white male candidates. In addition to bias, in the rush to deploy AI companies need to ensure they are following their other established ethical practices, including transparency and accountability.

There are significant copyright and intellectual property concerns around AI.

AI Correctness, Accuracy and Safety

Finally, “naïve” AI systems want to please and can sometimes generate false information. Generative AI constructs information. There have been some cases recently where AI systems “constructed” fake legal cases, which were submitted to the court. Additionally, AI systems can provide incorrect or unsafe advice. Recently an eating disorder website added a chatbot to answer questions, only to find later that the chatbot was suggesting to potentially anorexic website visitors that they should cut their daily intake by 500 to 1000 calories. It can be argued that these are more of an example of poor AI governance instead of any inherent failure of AI. A legal intern would never be allowed to submit his brief without proper review. The same processes need to be followed when using AI.

Despite these risks and concerns, IT and legal departments will face tremendous pressure in 2024 to deploy AI applications. Organizations may lose an advantage sitting on the sidelines. Waiting until the compliance and risk environment becomes better understood will not be an option for many.

Creating an AI Governance Program

Launching an AI-assisted application with risks and concerns or not using AI at all is a false dilemma. Today companies are successfully using AI compliantly, limiting legal risks, ethically and correctly through an AI Governance program. Good AI Governance not only drives effective program development, but it also saves time, allowing AI applications to move into production more quickly.




Figure 5. Key steps in developing and launching an AI Governance program.

STEP 1

Stakeholder Engagement and Assessment

Step 1: Engage Stakeholders and Conduct Assessments

While it may be tempting to try and develop a program with a small group of stakeholders, this may slow down or even halt program development. As a first step, needs should be assessed and socialized with a larger group of stakeholders early in the process. Tasks include:

Assessment and Roadmap development Assess current and targeted compliance and risk requirements, engaging key stakeholders in the process. Determining what needs to be done and at what level can speed up these types of complex projects.

Impact Assessment creation Some jurisdictions may specifically require an AI impact assessment focusing on people and organizations.

Organizing an AI Governance steering committee AI governance is complex, requiring the participation of various stakeholders, including compliance, risk, legal, privacy, information governance, data governance, IT, and business functions. This committee should be formed early in the process to ensure both that all risks and requirements are covered and that each group feels a sense of “buy in” to the process.

Roles and Responsibilities definition Many AI governance functions will require effort from groups. Establishing ongoing roles and responsibilities early ensures that AI governance will be an ongoing, continual process and not a one-time exercise.

STEP 2

AI Governance Policies

Step 2: Develop and Update Policies

Successful AI governance requires the creation of a new policy and possibly updates of existing policies.

Create AI Governance Policy creation The next step is developing and updating AI Governance policies. An AI Governance Policy sets out the compliant, transparent, and ethical use of AI for the organization. It details how AI should be used, safeguards employees, and ensures compliance with regulatory requirements. This is your overall “guiding light” to demonstrate to others you are using AI responsibly.

Data Retention/Records Retention Policy and Schedule update Organizations may need to update their data retention policies or records retention schedules to avoid older legacy data with sensitive or incorrect information “polluting” the development of AI systems.

Data Security Classification Policies update A data security classification policy classifies information based on privacy, confidentiality, intellectual property, and other sensitivity factors. It may need to be updated to ensure that appropriate controls are placed on sensitive information.

Privacy policies update Many AI regulatory restrictions center on the use of personal information. Privacy policies need to be synchronized with AI governance policies.

Having up-to-date policies provides defensibility if an AI system faces review from a regulator. These policies demonstrate the organization is mindful in its use of AI and is diligent in its compliance efforts.

AI Governance Process Development

Step 3: Develop AI Governance Processes

AI requires the development of governance processes. These processes will come into play both during development and during ongoing deployment.

Regulatory Review Process creation AI regulatory requirements are being announced seemingly every week. Organizations need to develop a process for monitoring regulatory changes to ensure their systems comply with any new rules.

Data Provenance Process development AI systems leverage both “training data” used by large language models and supplementary information used as part of retrieval augmented generation. Companies initially need to undertake reasonable due diligence to ensure this input data is not copyrighted or, if it is, that they have the right to use this information. Furthermore, as this input data is often refreshed, ascertaining provenance must occur periodically.

Privacy and Sensitive Information Review Process development In addition to ensuring that input data is not copyrighted, organizations should develop a process to ensure the AI does not contain either personal or other types of sensitive information such as trade secrets or corporate confidential information.

Ethical Use Review Process creation In addition to compliance, AI systems need to produce ethical results. For example, a visual generative AI application when asked to create a picture of “senior executives” should not consistently create an image exclusively consisting of older white males. The AI output should be evaluated to ensure it is being used ethically and reflects an organization’s values.

AI Accuracy, Correctness, and Safety Review Process development In addition to compliance, legal assuredness, and accuracy, AI needs to be accurate, correct, and safe. AI’s polished output can lull a user to believe that all the information it produces is correct and accurate. Correctness and accuracy need to be tested both throughout development and on an ongoing basis. Additionally, AI also needs to be tested for safety to ensure it is not being misused.

All AI Governance processes should be completed regularly, and the results should be retained.

AI Governance Process Execution, Monitoring, and Remediation

Step 4: Process Execution, Monitoring and Remediation

Once launched, AI systems need to be monitored. Any issues, discrepancies, or problems should be noted, along with steps taken to remediate these issues.

Initial and ongoing process execution Ensure all processes are enacted during and after launch.

Ongoing monitoring Monitoring and reviewing the results of the processes.

Updates and remediation Updating the system and/or approach as compliance and other rules change. Remediating any issues encountered, and documenting the actions taken.

In the event of a regulatory inquiry, being able to readily communicate what you intended to do (policies), how you intended to ensure you were doing it (processes), and how you addressed issues when they arose will demonstrate compliance and make the system more defensible.

Conclusion

A strong AI governance program that is developed early speeds up overall deployment. By identifying compliance requirements and anticipating risks during system development it can avoid having to redesign or rework the system on the tail end. Likewise, good governance engages key stakeholders early on, allowing them to both raise concerns as well as become comfortable with chosen approaches. While a bit counterintuitive, creating AI Governance drives faster development times.

Even after the program is launched, new regulations and other legal requirements will be enacted. How can organizations create an AI Governance function today that does not have to be updated every time a new rule is announced? A good initial program design can minimize the effort needed for future updates. Many of the global requirements address similar and common areas. Designing an AI governance functions toward these common requirements will enable companies to become “AI Agile,” requiring minimal changes.

This new, complex technology faces a chaotic legal and regulatory environment. However, good AI Governance follows established compliance and risk reduction strategies, even if applied in this near area. Fear not AI. Instead with a smart approach embrace the technology and profit from it.

About Contoural

Contoural is the largest independent records management, Information Governance, and privacy strategic consulting service with a focus on AI. Contoural does not sell any products or receive referral fees, serving as a trusted advisor to its clients by providing unbiased advice. Contoural has more than 30% of the Fortune 500 plus many midsize and small organizations as clients, across all industries, as well as federal agencies and local governments. Contoural offers a range of records management and Information Governance consulting services. Additional information is available at www.contoural.com or contact us at info@contoural.com.

- AI Impact Assessment
- AI Governance steering committee development
- AI Governance roles and responsibilities development
- AI Governance policy development
- Data and records retention policy update
- Data security classification policy update
- Privacy policy update
- Regulatory review process development
- Data provenance and assurance process creation
- Privacy and sensitive information review process creation
- Ethical use validation process creation
- AI Governance program updates and remediation

Disclaimer

Contoural provides information regarding business, compliance and legal trends and issues for educational and planning purposes. However, legal information is not the same as legal advice—the application of law to an individual's or organization's specific circumstances. Contoural and its consultants do not provide legal advice. Organizations should consult with competent legal counsel for professional assurance that our information, and any interpretation of it, is appropriate to each organization's particular situation.



335 Main Street, Suite B, Los Altos, CA 94022

650.390.0800 | info@contoural.com | www.contoural.com

© 2024 All rights reserved, Contoural. 011824