COSMOS
Navigate the Compliance Universe

CEP Magazine - March 2024

Mark Diamond (markdiamond@contoural.com) is the CEO of Contoural Inc. in Los Altos, California, USA.

Creating an AI governance function: Part 1

By Mark Diamond

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. Companies are launching AI initiatives in legal, finance, marketing, product design, engineering, and nearly every single aspect of the organization. (Full disclosure: our company, Contoural, is launching an AI-based records management initiative.) Without overstating, AI has the potential to be transformative. Part 1 of this two-part series addresses the risks and restrictions organizations face in deploying AI and the key elements of an AI governance strategy. Part 2 will detail how to develop an AI governance function.

Regulations are evolving

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.[1]

In the U.S., states are limiting how companies can use AI to make financial decisions such as loan approvals. (Note: 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.[2] Recently, the U.S. Securities and Exchange Commission (SEC) warned businesses against "AI washing," or making false claims about their use of AI.[3] Companies are feeling so much pressure to show their investors that they are taking advantage of AI-powered products; therefore, the SEC felt it necessary to caution against making bogus claims.

These new regulations are just the beginning, as we expect to see many countries and states developing new rules limiting AI this year. The forecast for the AI regulatory environment is both rushed and a bit messy.

Emerging AI regulatory requirements

Governments across the world are quickly enacting a variety of AI regulations. Here is a small sampling of different requirements.

Europe: Europe considers itself a leader in AI regulation and has created a legal, regulatory framework largely based on potential risk posed by AI systems. High-risk systems would be more strictly regulated, including being required to carry out a rights impact assessment. Lower-risk systems would have fewer regulations, including disclosure requirements.

U.S. federal government: The Biden administration issued an executive order on AI. It requires that AI developers share results of safety tests with regulators, provide guidelines for the federal government's own use of AI, and prohibit AI-driven discrimination.

U.S. states: Numerous states have either proposed or enacted a variety of AI regulations focusing on disclosure, consumer profiling, unfair discrimination in financial services and hiring, facial recognition, and registries.

ISO24001: The International Organization for Standardization recently released a framework for organizations involved in developing, providing, or using AI-based products or services.

China: China has developed a series of AI regulations, including a government registry and rules limiting synthetically generated images, video, audio, and text.

These requirements—and the number of jurisdictions creating them—are certain to grow in the coming year.

Sitting on the sidelines is also a risk

Aside from regulatory requirements, AI raises a host of legal issues around copyright protections, intellectual property ownership, product liability, labor and employment, and other areas. 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.

Yet, despite the ever-changing legal and regulatory environment, companies are facing tremendous pressure to develop and deploy AI-based applications. The investment community sees AI's potential and will reward companies with increased share prices for those perceived to be taking advantage of it. Senior management will feel pressure to show that the company is using AI to move its share price.

Thus, 2024 is setting up for a tremendous tug-of-war between the parts of the business that want to use AI and the chaotic compliance and risk environment of using it. Organizations may lose an advantage sitting on the sidelines, waiting until the compliance and risk environment becomes better understood.

What should companies do? They should start creating AI governance functions today to develop policies, processes, and organizational structures to meet compliance requirements and ensure AI's legal, ethical, and accurate use. While generative AI may be new, companies can apply established compliance and risk management strategies to efficiently navigate these challenges and enable the successful use of these technologies.

Key AI terms

Generative AI: A form of AI capable of generating text, images, or other media using generative models that learn the patterns and then put out information with similar characteristics.

Large language models (LLMs): AI "engines" that transform large amounts of information to understand, summarize, generate, and predict new content. LLMs are either proprietary to a specific AI vendor (closed) from companies like Open.AI, Google, or Anthropic or available through open source (open model) from companies like Meta.

Hallucinations: When generative AI creates answers that are false but presented as real or factual, such as citing a nonexistent legal case as support for a briefing.

Training data: Developers feed large amounts of data into AI model learning algorithms to make decisions.

Retrieval augmented generation (RAG): A technique in which the user supplements training data with their own data. In developing a contract generation system, for example, RAG would allow a user to provide samples of their own, more relevant contracts into the system.

Data provenance: Understanding the origin of a piece of data in a database, document, or repository, together with an explanation of how and why it got to the present place. Data provenance is important to ensure any training data copyright and intellectual protections are respected.

AI safety: Preventing accidents, misuse, or other harmful consequences of AI. AI safety should prevent a large language model from answering a query, for example, on how to make illegal explosives, even if the system has the knowledge to answer the question.

Key elements of AI governance

Many recent media reports on AI have focused on specific impacts, including privacy and accuracy. While these are significant, AI impacts many other additional risks—and organizations must address all of them. Key elements of AI governance include the following.

Compliance

Many governments have announced their intention to regulate the use of AI. Current regulatory efforts are somewhat scattershot, loosely falling into the following categories:

- **Privacy:** Restrictions on the use of personal information to be used as input for AI-automated decision-making and classification.
- **Nondiscrimination:** Preventing discrimination by AI on race, age, or other factors in offering or denying products or services. For example, AI systems must not offer lower credit limits to women than men.
- **Transparency:** Requirements for AI systems to process information in a transparent manner. For example, if an AI system hallucinates, the system should be able to identify what caused the system to hallucinate.

Legal assurance

Designers of AI systems and applications must be careful not to infringe on copyrighted material or other intellectual property to which they do not have rights. The concern is twofold: (1) the designers do not have the authority to use the copyrighted material for the purpose of training their AI (that is, that it is not a "fair use" of the copyrighted material), and (2) if training data used within a large language model is copyrighted material, the output itself may infringe. Recently, the New York Times sued OpenAI and Microsoft alleging their AI products infringe on their copyrighted articles.[4] Often, the user may not know the data provenance of the training data in closed or proprietary large language models and must depend on the assurances of the model's vendor.

Sensitive information

In addition to ensuring AI does not infringe on copyrighted content, AI deployers must ensure they are not inadvertently supplying the AI model with sensitive information, including trade secrets, corporate confidential, or personal information.

Ethical use

AI systems must drive ethical decisions. For example, if using AI to screen potential job applicants, measures need to be taken to ensure the AI does not inadvertently disfavor diverse applicants or create other inadvertent negative impacts (sometimes referred to as "legal or other significant effect").

Accuracy and correctness

AI systems must produce factually correct and accurate results. This includes avoiding hallucinations, in which the system fabricates information presented as fact or provides misinformation.

Safety

AI safety ensures that systems are not misused for unintended or nefarious purposes. For example, an AI system that creates chemical processes should not be able to describe how to create illegal explosives.

Clearly, there is a fair amount of crossover among the categories in this framework. An AI system that unethically discriminates against job candidates would be in violation of employment rules and risk legal liability.

Creating AI governance agility

While there are a multitude of AI rules, many of these focus on a common set of requirements or restrictions. AI governance functions should not be geared toward specific rules but rather focus on common areas across requirements. Creating "AI governance agility" will meet many of the specific requirements found across the world. This approach will be far easier in the long run than developing a hardwired program for a single jurisdiction, only to constantly update the program as jurisdictions adopt new legislation.

Conclusion

Compliantly, defensibly, ethically, and correctly using AI poses real challenges for organizations. As regulators race to impose new rules and the legal intellectual property issues are being addressed by the courts, some compliance professionals may hope to delay the launch of AI initiatives until there is better clarity and these issues are settled. However, AI's productivity and economic benefits are too great for organizations to wait. Rather, forward-thinking compliance professionals will take up these challenges today.

Part 2 of this two-part series details how organizations can develop an AI governance function. An effective AI governance function is a combination of engaging the right stakeholders, developing and updating policies, generating a series of processes, and executing the processes consistently. While generative AI may be known, developing effective AI governance relies on applying proven compliance strategies. You know more than you think.

Takeaways

- Organizations are facing a compliance tug-of-war—pressured to deploy artificial intelligence (AI) despite significant compliance, legal, and other risks, as well as operational challenges.
- Companies are feeling pressure to show their investors how they are using AI.
- Regulators around the world are proposing and enacting AI regulatory requirements.
- AI systems' use of copyrighted information as an input creates legal risk.
- A comprehensive AI governance strategy addresses compliance, legal assurance, sensitive information,

ethical use, accuracy, compliance, and safety.

1 European Parliament, "EU AI Act: first regulation on artificial intelligence," news release, updated December 19, 2023, https://www.europarl.europa.eu/news/en/headlines/society/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence.
2 The White House, "FACT SHEET: President Biden Issues Executive Order on Safe, Secure, and Trustworthy Artificial Intelligence," news release, October 30, 2023, https://www.whitehouse.gov/briefing-room/statements-releases/2023/10/30/fact-sheet-president-biden-issues-executive-order-on-safe-secure-and-trustworthy-artificial-intelligence/.
3 Richard Vanderford, "SEC Head Warns Against 'AI Washing,' the High-Tech Version of 'Greenwashing,'" *The Wall Street Journal*, December 5, 2023, https://www.wsj.com/articles/sec-head-warns-against-ai-washing-the-high-tech-version-of-greenwashing-6ff60da9.
4 Haleluya Hadero and David Bauder, "The New York Times sues OpenAI and Microsoft for using its stories to train chatbots," *Associated Press*, December 27, 2023, https://apnews.com/article/nyt-new-york-times-openai-microsoft-6ea53a8ad3efa06ee4643b697df0ba57.