



WHITE PAPER | OCTOBER 2024

SAGE: A systematic approach to data-driven AI governance

SAGE: A systematic approach to data-driven AI governance

Authors: Danny Tobey, M.D., J.D., Ashley Carr, J.D., Karley Buckley, J.D., Kyle Kloeppel, J.D., Sam Tyner-Monroe, Ph.D.

Introduction

Artificial intelligence (AI) technology is advancing at an unprecedented rate, increasing in complexity while driving significant innovation across sectors. In response to the rapid development of AI, myriad compliance frameworks are being established to ensure the responsible development, deployment, and use of this emerging technology.

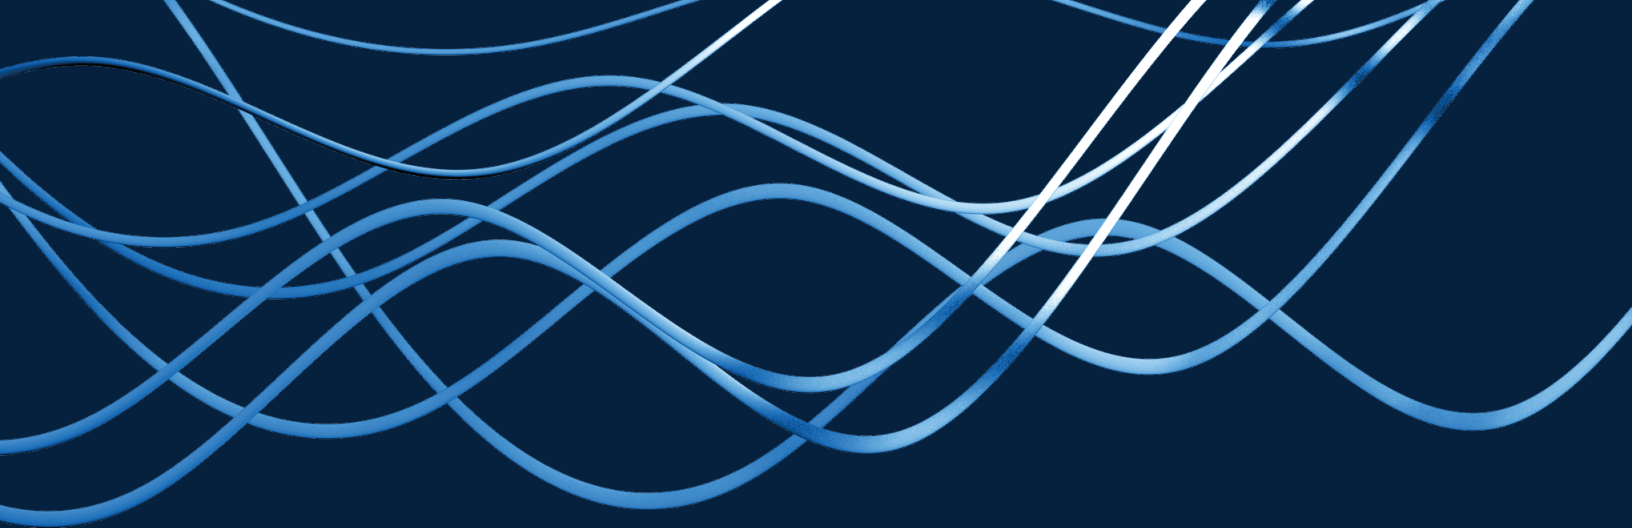
These frameworks include regulatory schemes, newly enacted laws, and a growing body of industry and technical standards, each with the goal of furthering safety, trustworthiness, and fairness. Providers, developers, and deployers of AI systems are acutely aware of the need to implement robust controls around their interactions with AI – not only to comply with the law, but also to maintain public trust, mitigate litigation risk, and ensure they are using AI responsibly.

Unfortunately, implementing the necessary controls is far from straightforward. On the surface, there is an overabundance problem. Regardless of one's philosophy on regulating AI, as a practical matter, many organizations are simply unable to keep up with the sheer volume of requirements being promulgated, let alone implement controls to address each and every

one. Already, global, national, state, and even local AI regimes are creating a web of intensive requirements. Moreover, horizontal "omnibus" regulation, governing AI per se, is overlapping with vertical regulation within industries and domains, like insurance and healthcare.

Beyond the growing "framework fatigue," the individual regulations comprising this patchwork are each often remarkably complex, rivaling the complexity of the technology itself. And, despite the length and density of these efforts (the European Union Artificial Intelligence Act, or EU AI Act, alone is 458 pages), much remains unknown. This causes a secondary convergence problem: It is often difficult to understand the nuances of these numerous requirements, how to make them actionable, and how to address the gaps and ambiguities – not to mention how these individual regulations interact with each other.

And yet, companies face potentially massive penalties for AI compliance failures (under the EU AI Act, up to seven percent of global annual revenue; under the FTC, fines and even loss of expensive algorithms and data). Meanwhile, boards and C-suites are demanding rapid adoption.



Addressing this problem has called for substantial effort, as well as technical and legal contributions combined. Having advised numerous companies, including several of the Fortune 10, on AI governance, compliance, and regulation, our key to solving these problems has been the creation of a systematic approach: the Standardized Assessment and Governance Enhancement (SAGE) framework.

This approach entails combining levels of human and machine synthesis and analysis into a consolidated and grounded guiding source, from which organizations can understand the contours of emerging AI requirements and their interactions with each other, assess compliance of their existing internal systems and processes, and implement a set of risk-based controls that address apparent gaps and are fit to purpose and risk profile.

In this paper, we will begin in Section I with an overview of the compliance landscape, including describing in detail both the problems of overabundance and convergence. In Section II, we will describe our method of building a comprehensive methodology to address these problems. Then, in Section III, we will describe

how we leveraged this novel methodology to provide companies with actionable guidance to understand and fulfill their obligations under the developing compliance landscape.

Given the sheer complexity of the problem, perfection is impractical (indeed, given conflicting standards, likely impossible). What we propose is a reasoned approach, drawing thoughtfully from the many frameworks and regulations, with documentation to explain the “how” and the “why,” if and when regulators ask.



Section I: The complexity of AI compliance obligations

In this section, we discuss two major problems faced by organizations during AI compliance efforts: tracking the overabundance of new laws and standards and understanding how the requirements of each converge and interact with one another.

The overabundance problem

The rapid advancement of AI has caught many by surprise, including the public, lawmakers, and regulators. The sophistication and capability of AI systems – particularly foundation models – have seemingly overnight demonstrated a level of power that was previously resigned to science fiction. These technologies are performing tasks that were once considered the exclusive domain of human intelligence. The sudden realization of AI's potential and its attendant risks if left unchecked has led to widespread recognition of the need for oversight.

The recognition that oversight is necessary has driven a significant push to regulate the development, deployment, and use of AI. The seemingly abrupt leap in AI capabilities has highlighted the dangers of operating without a robust regulatory framework, including ethical concerns, security risks, and the potential for unintended, unforeseen consequences.

President Joe Biden stated in the October 2023 Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence that “artificial intelligence holds extraordinary potential for both promise and peril” and that the “interests of Americans who increasingly use, interact with, or purchase AI and AI-enabled products in their daily lives must be protected.”¹ Lawmakers, regulators, and standards-setting organizations around the world share this view and are simultaneously working to create a legal and regulatory framework meant to address AI risks.

As a result of this movement to regulate AI, an unwieldy number of AI specific laws, regulations, and industry standards have been introduced across many jurisdictions. Legislative mentions of AI in global proceedings rose from 1,247 in 2022 to 2,175 in 2023.² Between 2016 and 2023, more than 128 countries proposed at least one AI-related law, with 32 of those countries enacting at least one AI-related law.³ In the US alone, 181 AI-related laws were proposed at the federal level in 2023.⁴ At the state level, 45 US states and

1. Executive Order 14110, 88 FR 75191 (2023).

2. Nestor Maslej, Loredana Fattorini, Raymond Perrault, Vanessa Parli, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, Juan Carlos Nieves, Yoav Shoham, Russell Wald, and Jack Clark, “The AI Index 2024 Annual Report,” AI Index Steering Committee, Institute for Human-Centered AI, Stanford University (2024).

3. *Id.*

4. *Id.*

territories introduced AI-related bills in 2024, and 35 states enacted laws or adopted resolutions.⁵

Beyond the sheer number of proposed laws and regulations, a complicating factor is that this regulatory activity is happening at the federal, state, and local levels of government. What's more, existing laws and regulations are also being amended to account for AI's effects on their scope. This leaves organizations in a state of information overload. Developers and deployers of AI must track hundreds (if not thousands) of emerging requirements in order to even begin defining their compliance obligations. This volume of information is not digestible or actionable and is itself a bottleneck to compliance.

The convergence problem

The overabundance of emerging AI requirements leads to a secondary and more substantive problem – determining whether, when, and how to comply with them all. Organizations recognize the need to build AI risk management and compliance frameworks, but the sheer volume of requirements being developed by regulators, lawmakers, and standard-setting bodies is not actionable in its raw form. Organizations are therefore faced with the task of distilling and synthesizing these complex, overlapping, and even conflicting requirements.

Organizations must understand the areas of overlap, which requirements apply and when, what exceptions apply, and how to develop a compliance program that meets those requirements – on top of their existing technology-agnostic and sector-specific compliance obligations. Failing to adhere to applicable laws and regulations can have serious consequences for organizations.

For instance, penalties for noncompliance with certain provisions of the EU AI Act include fines up to the greater of EUR35 million or seven percent of an organization's global annual turnover,⁶ violations of Colorado's AI Act may result in injunctive relief or fines up to USD20,000 per violation,⁷ and federal regulators are consistently pursuing algorithmic and data disgorgement as a remedy of choice in AI regulatory enforcement actions.⁸

And even beyond AI-specific laws, AI developers and deployers are facing an onslaught of enforcement activity applying old laws to new technologies. As the FTC has put it, "There is no AI exemption from the laws on the books,"⁹ with historical fines in the millions and billions. Add to that an onslaught of civil litigation ranging from intellectual property infringement claims to consumer protection class actions to novel tort claims. Needless to say, the stakes are high when it comes to meeting organizational compliance obligations, but the path to get there is challenging and uncertain.

5. National Conference of State Legislatures, Artificial Intelligence 2024 Legislation (September 24, 2024), <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2024-legislation>.

6. Regulation (EU) No. 2024/1689, https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L_202401689.

7. Colo. Rev. Stat. Ann. § 6-1-112.

8. See, e.g., Press Release, Federal Trade Commission, Rite Aid Banned from Using AI Facial Recognition After FTC Says Retailer Deployed Technology without Reasonable Safeguards (December 19, 2023), <https://www.ftc.gov/news-events/news/press-releases/2023/12/rite-aid-banned-using-ai-facial-recognition-after-ftc-says-retailer-deployed-technology-without>.

9. See, e.g., Press Release, Federal Trade Commission, FTC Announces Crackdown on Deceptive AI Claims and Schemes (September 25, 2024), <https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes>.

Section II: Building the SAGE framework

Given the challenge of obtaining reasonable compliance with multiple overlapping authorities and rules, we have devised a methodology for companies at heightened risk – whether by size, prominence, industry, consumer orientation, or use cases – to assess, obtain, and document compliance with AI standards. The scale of the problem begs a technical solution, while the organizational and regulatory complexity and nuance require legal experience and customization.

We have achieved that balance through a technology-enabled legal approach, with maturity assessment, program build, validation, and documentation accomplished through a complementary and iterative human and AI analysis. This approach enables organizations to make an informed, risk-based decision about where and how to focus compliance efforts.

In our work, we have created a qualitative and quantitative central guiding source to ground organizational compliance efforts. Against such a source, organizations can work to accomplish three objectives:

- 1. Understand the intricacies of compliance obligations.** With a central guiding source, organizations can understand their compliance obligations as well as how each emergent or existing law, regulation, and standard overlaps with one another.
- 2. Assess their current governance programs against a standard framework.** A single guiding source allows organizations to assess and measure their compliance with applicable requirements, as well as understand their organizational maturity vis-à-vis these requirements.
- 3. Build compliant and rightsized AI governance programs, mapped to regulatory standards.** Using a consolidated source as a guide, companies are able to build compliance programs that are rightsized to their organization, and at the same time, adhere and address existing requirements.

DLA Piper’s Artificial Intelligence and Data Analytics practice has created the SAGE framework as this consolidated point of reference for cross-regulatory program building. To create the SAGE framework, we leveraged both the power of AI and the human legal and technical experience of our team. Moreover, using this guiding source, we have created methods and work product to pursue the three objectives outlined above and help guide organizations in their pursuit of compliance.

Database creation methodology

Creating the SAGE framework begins with a process we call “atomization.” Generally speaking, this process consists of taking a composite product such as a regulation and breaking it into its component parts. We applied the atomization process to leading AI related laws, regulations, and standards.

To accomplish this, we deployed our attorneys to analyze sources such as the EU AI Act, the National Institute of Standards and Technology AI Risk Management Framework, and the International Organization for Standardization 42001 standard, among others. Based on their analysis, these attorneys then extracted thousands of discrete requirements contained within these source documents and stored them in a database.

This manual atomization approach leverages attorneys’ ability to interpret and analyze complex legal text and requirements. The complexity, nuance, and context required for accurate legal interpretation throughout the atomization process often exceeds current AI and technical capabilities.

As an example of this process, take, for instance, a standard that requires organizations to “develop an AI policy that allocates responsibility for AI oversight

within the organization.” Through the atomization process, this would become the following: (1) The organization shall develop an AI policy and (2) the organization’s AI policy shall allocate responsibility for AI oversight within the organization. By engaging in this process, we were able to capture and understand the extent of the actions and controls required by each source.

In creating this database, DLA Piper attorneys and data scientists devised a method intended to both eliminate any unnecessary duplication within the source documents analyzed as well as memorialize overlapping atomized requirements across sources.

Eliminating duplication

The atomization process for any one source often results in many hundreds of individual requirements, some of which may be internally duplicative. In order to ensure internal consistency of the atomic requirements within each source document and consolidate the number of requirements atomized, we use a two-step process that leverages natural language processing (NLP) techniques and attorney experience.

First, we used an embedding algorithm to understand the similarity between each atomized requirement. In order to do this, our chosen model compares each atomized requirement against every other atomized requirement within the same source and indicates how semantically similar the requirements are to each other.

The algorithm creates a numerical representation of each requirement’s meaning based on hundreds of characteristics, known as dimensions. Using these numerical representations, the algorithm then measures the similarity between the requirements, which is determined by how close two requirements are to one another in a high-dimensional

space. This process results in a similarity score between every possible pair of requirements, which quantifies the semantic similarity between each pair.

Next, attorneys review pairs of atomized requirements with high similarity scores and consolidate any duplicative criteria. This process allows us to ensure that large standards, laws, or regulations are internally consistent. As a result, we are left with a consolidated atomized standard that reduces repetition and applies the fewest number of unique requirements possible.

Memorializing overlap

A primary goal of our atomization process is not only to ensure that we are left with a nonduplicative set of requirements extracted from each source, but also to understand how each of the atomized requirements overlap with one another across sources. In order to understand and record this overlap, we leveraged a similar two-step process as described above, using both NLP techniques as well as experienced attorneys.

First, as new sources are atomized and consolidated, we use the same methods described above to compare each new atomized requirement against the existing atomized requirements in our current database. Attorneys then review pairs of atomized requirements with a high similarity score and consolidate duplicative requirements, noting the sources containing the consolidated requirement. As a result of this process, we are able to understand and record the overlap of atomized requirements across sources.

Leveraging the processes described above and depicted below in *Figure 1*, we created a database that contains a comprehensive set of atomized requirements from leading sources, which also limits redundancy and captures how these sources overlap.

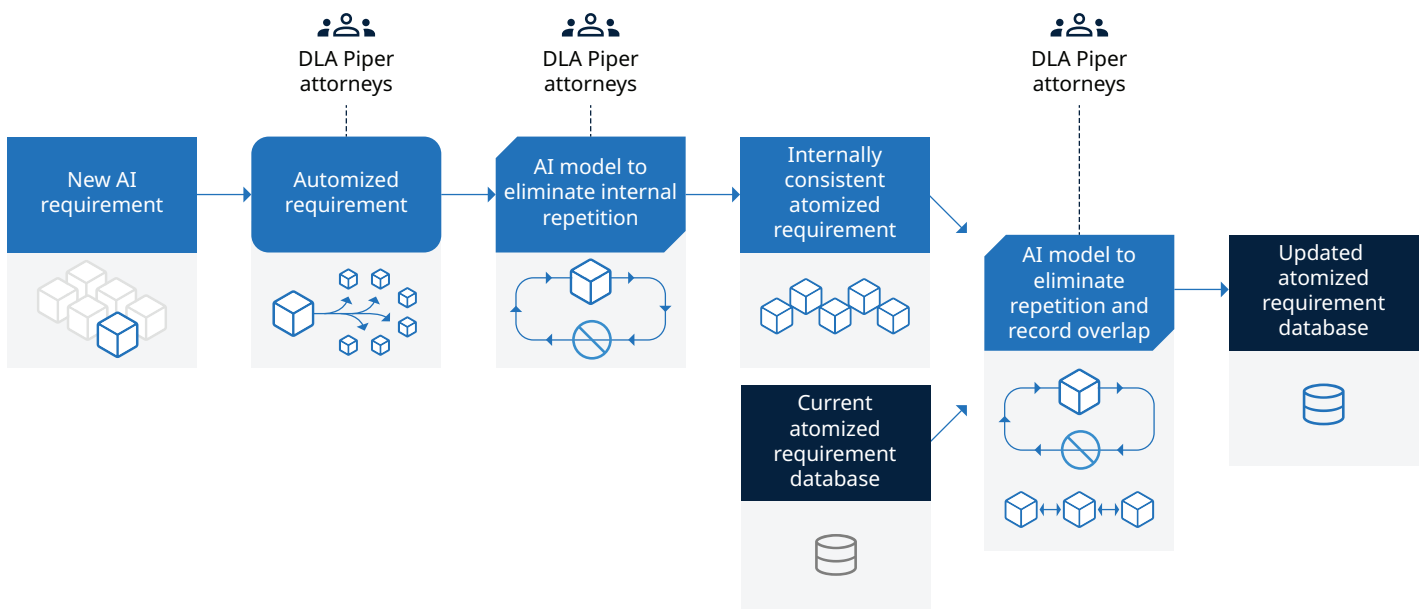


Figure 1: The SAGE process for consolidating many global AI requirements into a cohesive set of atomized requirements.



Section III: Creating actionable solutions

Using our novel SAGE framework, DLA Piper attorneys are able to provide our clients with actionable guidance to address their compliance needs. This guidance includes providing grounded, cross-jurisdictional advice and resources related to existing compliance requirements, assessing organizational governance programs against those requirements, and constructing thoughtful and compliant AI governance programs.

Crosswalk creation methodology

Our clients know the importance of staying abreast of industry standards and legislation in the AI space, but very often have questions about which standard(s) to adopt, whether compliance with one will necessarily result in compliance with another, and whether there are material differences in compliance obligations that will impact their business.

In our experience, organizations navigating an uncertain legal environment can benefit from actionable, risk-based guidance. With those needs in mind, our next step was to use the atomized database to distill and “crosswalk” (*ie*, map) the core requirements within and across the leading standards. The crosswalk enables organizations to quickly understand how each standard addresses a given topic (*eg*, impact assessments or AI system testing) and begin to align on required organizational controls.

To create the crosswalk, we started with the raw data in the atomized database. First, we used semantic clustering models to group together similar atomized criteria. Using AI to take the first cut at sorting this large amount of raw material enabled our team to more efficiently identify responsible AI dimensions covered by each standard (what we call “compliance categories”). Then, the semantically similar groupings were analyzed and synthesized by our AI Governance team of attorneys.

Using their domain experience regarding risk mitigation and compliance strategies, the attorney team distilled the key requirements for each compliance category. Through this process, thousands of atomized criteria across seven leading standards and regulations were distilled into a crosswalk that maps the key requirements across two dozen compliance categories ranging from AI inventory to bias mitigation to AI system decommissioning.

The purpose of the crosswalk is to give organizations and legal teams a mapping of key attributes of the myriad of emerging AI laws and standards at three levels of granularity. While the atomized database is encyclopedic, the crosswalk extracts actionable distillations.

First, the crosswalk contains a single-page matrix that allows a reader to quickly visualize the compliance categories covered by each standard and how they may overlap. Then, in the second section, it provides a distilled summary of the key, high-level takeaway of what each standard requires across the compliance categories. Finally, the third section is a detailed mapping of key requirements with citations to key provisions. This map facilitates harmonization of requirements across standards, which in our experience is core to developing cross-jurisdictional, future-proofed AI governance and compliance programs.

Evaluation methodology

The second practical application of our atomized database is for enhancement of lawyer-led maturity assessments to identify gaps and areas of strength in an organizations' existing AI governance and compliance program. The atomized criteria form the backbone of a standards-backed maturity assessment in several ways.

First, the criteria inform the categories of documentation and information that we request from an organization in order to perform this assessment. Second, we are able to use generative AI with lawyers-in-the-loop to assess the client's existing documentation against the database of atomized requirements. Then, finally, we are able to use the results of the AI-enabled assessment to evaluate specific compliance gaps and assess their materiality in context, leading to an evaluation of the organization's overall AI governance maturity. This maturity evaluation delivery process is depicted in *Figure 2* below.

By using the atomized database as a rubric against which an organization's existing controls can be assessed, we achieve a more comprehensive, standards-backed maturity assessment that is both qualitative and quantitative. While the human attorneys perform a standard qualitative, experience-based evaluation, they also run the computational mapping exercise, evaluate the "heat map" of potential gaps against various applicable standards, and then apply human legal reasoning to that quantitative analysis. Combining the quantitative and qualitative approaches provides a robust maturity and gap assessment with the two modalities complementing one another.

Rightsized AI governance

We use our atomized criteria to inform the governance programs we design for our clients. This tool enables our team to tie the content of each instrument we create to specific regulatory or industry requirements, supporting a comprehensive compliance program. Paired with the evaluation methodology outlined above, we are able to first pinpoint gaps and then design controls that are specifically tied to those requirements. Human-centered legal analysis interprets and applies these findings with judgment, ensuring that governance is rightsized to an organization's AI risk profile, business needs, and compliance culture, and helps to prepare for regulatory scrutiny by documenting a systematic and cross-referenced approach for each compliance category.

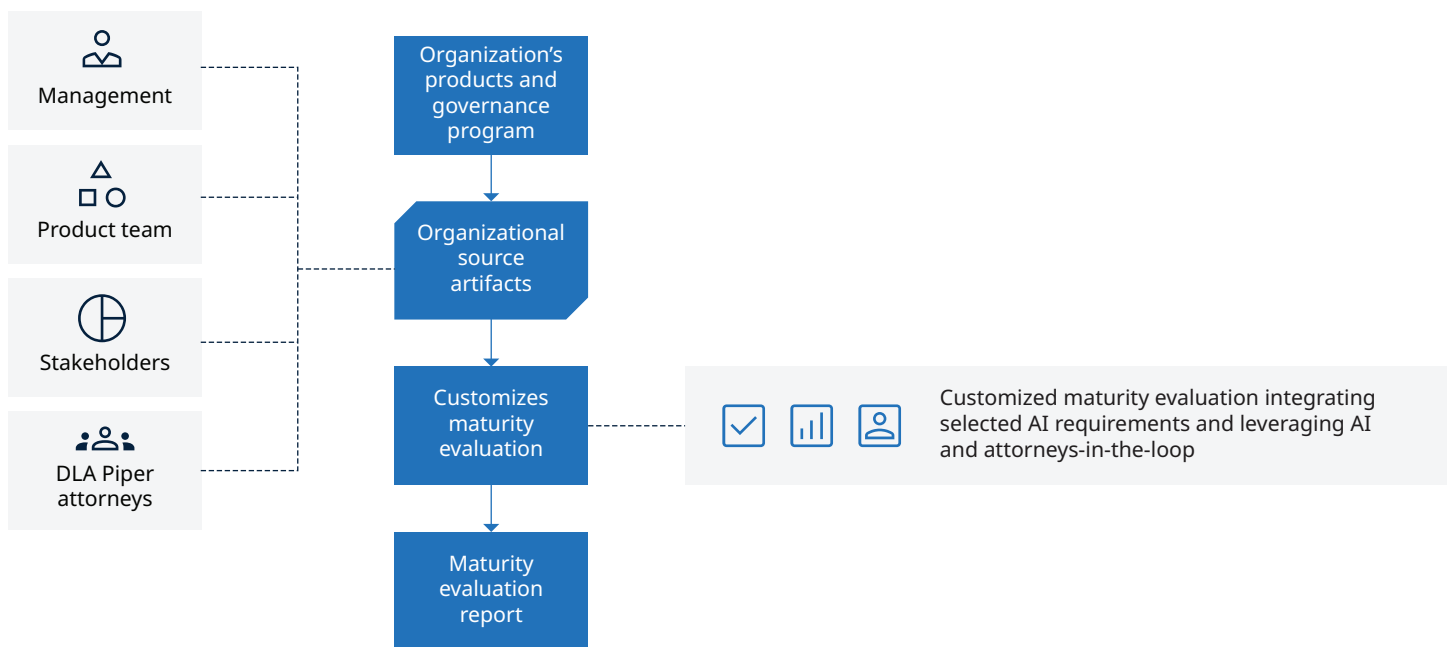


Figure 2: The evaluation delivery process for assessing organizational maturity against AI requirements.

Conclusion

Our team has presented a novel and comprehensive method of mapping AI regulations and standards and creating standards-backed AI governance programs in a data-driven manner. We have explained how we employ a hybrid approach that leverages legal subject matter knowledge and AI analysis to better deliver actionable and standards-backed governance and compliance programs.

We hope that this paper will encourage organizations to approach AI governance in a systematic and data-driven way, and inspire legal practitioners to continue to explore ways in which we can leverage emerging technology to deliver the highest-value legal services to clients.

About us

DLA Piper is a global law firm with lawyers located in more than 40 countries throughout the Americas, Europe, the Middle East, Africa, and Asia Pacific, positioning us to help companies with their legal needs around the world.

For more information

To learn more about DLA Piper, visit dlapiper.com or contact:



Danny Tobey, M.D., J.D.

Partner and Global Co-Chair and Chair of DLA Piper Americas AI and Data Analytics

T +1 214 743 4538

danny.tobey@us.dlapiper.com



Ashley Carr, J.D.

Partner and AI Governance Lead

T +1 817 713 5113

ashley.carr@us.dlapiper.com



Karley Buckley, J.D.

Associate, AI Governance

T +1 713 425 8421

karley.buckley@us.dlapiper.com



Kyle Kloeppe, J.D.

Associate

T +1 916 930 3253

kyle.kloeppe@us.dlapiper.com



Sam Tyner-Monroe, Ph.D.

Managing Director, Responsible AI

T +1 202 799 4522

sam.tyner-monroe@us.dlapiper.com



Innovative Practitioner:

Danny Tobey

Financial Times 2023

Top AI Lawyer: Danny Tobey

Insider 2022

Innovative Lawyers in

Technology

Financial Times 2023

Best Use of AI

Law.com 2024