Building a Comprehensive AI Governance Framework in Life Sciences

Gary F. Giampetruzzi/Amir R. Ghavi/Ann E. Beasley/ BJ D'Avella/Sarah Volden/Jessica Montes/ Julie Kudyba/Ravneet Talwar

This is the first in a series of articles by the Paul Hastings Life Sciences and Healthcare practice. In Part 1, the authors discuss the importance and regulatory drivers of implementing an Artificial Intelligence (AI) Governance Framework and how pharmaceutical and medical device companies can structure their governance program.

Gary F. Giampetruzzi is a partner in the Litigation Department and Chair of the Life Sciences & Healthcare Practice at Paul Hastings LLP.

Amir R. Ghavi is a partner in the Corporate Department of Paul Hastings LLP and co-chair of the firm's Technology Transactions group.

Ann E. Beasley is a Managing Director in the Life Sciences & Healthcare Consulting Group at Paul Hastings LLP.

BJ D'Avella is the Group Leader of the Life Sciences & Healthcare Consulting Group at Paul Hastings LLP.

Sarah Volden is a member of the Life Sciences & Healthcare Consulting Group at Paul Hastings LLP.

Jessica Montes is of counsel in the Litigation Department at Paul Hastings LLP.

Julie Kudyba is the Chief Ethics & Compliance Officer at Bausch Health.

Ravneet Talwar is a member of the Life Sciences & Healthcare Consulting Group at Paul Hastings LLP.

This article was originally published on Paul Hastings LLP's website and is reprinted with permission.

Pharmaceutical and medical device companies (collectively, life sciences companies) are organizations that impact the health, safety, and well-being of patients. These companies operate in a heavily regulated environment that requires a strong corporate focus on product quality, patient safety, and transparency. Long treatment development life cycles, which may span decades, further complicate the environment. Together, these and other factors imply a complex compliance framework that requires careful and thoughtful navigation when implementing novel operations and strategies involving AI applications. Life sciences companies are rapidly adopting such applications, which are transforming areas such as drug discovery and research and development. As examples, global pharmaceutical companies have entered into partnerships with AI tech companies to accelerate drug development in therapeutic areas of unmet need and have developed internal AI-powered platforms to streamline clinical trial processes. There are a variety of stakeholders excited by the unparalleled economies associated with integrating AI and generative AI technologies to dramatically improve patient outcomes.[1]

> *The term "artificial intelligence" means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments.*[2]

As AI adoption accelerates across life sciences, companies are building the structures needed to scale these technologies responsibly. Although the industry is known for its deep scientific expertise and innovation, AI adoption represents a shift toward a different kind of tool that challenges traditional models of discovery and demands new ways of working. The pace of AI advancement stands in contrast to the industry's historic development timelines, and while this has contributed to a more measured adoption curve, companies increasingly recognize the risk of falling behind in a landscape that is being rapidly reshaped by AI. To ensure a solid foundation for future AI-related activities, organizations should implement a clear, consistent, and comprehensive AI governance framework that allows them to realize the benefits of AI while effectively identifying and mitigating the risks associated with the rapidly evolving AI global regulatory landscape.

> ***"The first step in AI governance should be to ensure that current regulations apply to AI, to the greatest extent possible. If human activity without the use of AI is regulated, then the use of AI should similarly be regulated."***[3]

Globally, laws are being enacted, and regulations and guidance are being issued to address AI. These efforts generally seek to ensure the safety, security, and trustworthiness of AI systems while also fostering and supporting AI innovation. However, approaches to regulating AI varies from jurisdiction to jurisdiction.

For example, in the European Union, regulators are implementing a comprehensive, risk-based regulation, the EU AI Act, that applies to all "AI systems" developed, distributed or used in the EU. It establishes different rules according to the system's level of risk to health, safety and fundamental rights and imposes stringent requirements on high-risk systems (e.g., establish a continuous risk management system; train AI models on data sets that meet quality criteria).[4] Importantly for life sciences companies, the law classifies as high risk AI systems that are a safety component of or constitute a medical device under certain EU medical device and in vitro diagnostic medical device regulations (e.g., devices that are used for diagnosis, monitoring, prevention, prediction and similar uses) and are required to undergo a third-party conformity assessment. Thus, life sciences companies involved in developing, deploying, or using this type of AI will be required to adhere to the EU AI Act's most stringent requirements.

In the United States, there is no overarching federal law and the only state to have passed a comprehensive, risk-based law is Colorado, which postponed its implementation from February 1, 2026, to June 30, 2026, amid concerns about complexity and burden. If the Colorado law takes effect, certain AI systems used in healthcare—those that make or are a substantial factor in consequential decisions affecting healthcare services—will be considered high risk and therefore subject to the law's governance and disclosure requirements. Other states, such as California, Colorado, Texas, and Utah have passed far less comprehensive statutes aimed at specific applications (e.g., Utah's mental health chatbot law) or AI issues (e.g., California's training data transparency law).

At the federal level, policy has shifted from being pro-regulation and enforcement under the Biden administration to being firmly against regulation that stifles innovation. On July 23, President Trump released "America's AI Action Plan" (the Action Plan), which is structured around three core pillars: Accelerating AI Innovation, Building American AI Infrastructure, and Strengthening U.S. Leadership in International AI Diplomacy and Security. The Action Plan's pillar for "Accelerating AI Innovation" criticizes "onerous regulatory regime[s]" and includes recommendations that federal agencies "remove red tape" by revising or repealing regulations, rules, and other materials that "unnecessarily hinder AI development or deployment." The pillar also recommends that agencies take actions to "enable AI adoption" in "critical sectors, such as healthcare," including by establishing regulatory sandboxes at the Food and Drug Administration (the FDA). This Action Plan advances President Trump's January 23, 2025, "Executive Order on Removing Barriers to American Leadership in Artificial Intelligence," which revoked an AI-related executive order under the prior administration perceived as imposing "barriers to American AI Innovation."[5] The January 2025 executive order announced the current administration's policy of sustaining and enhancing American dominance in AI.

Federal agencies have removed certain Biden-era policies and guidance documents from their websites, and regulators are still developing frameworks to implement the AI Action Plan.[6] For now, life sciences companies should note that several Biden-era guidance documents and regulations within the Department of Health and Human Services remain in effect—which relate to specific applications and implicate patient safety. For instance, the FDA has maintained draft and final guidance documents issued in 2024 and January 2025 regarding marketing submissions for AI-enabled medical device software.[7] Additionally, the Office of the National Coordinator for Health Information Technology has maintained its 2024 update to the Health Information Technology Certification Program to require certain AI disclosures and risk management in health information technology, such as electronic health records systems.[8]

Life sciences companies should also note that the U.S. Department of Justice's (the DOJ) "Evaluation of Compliance Program Guidance" continues to contain key questions about the extent to which companies manage AI-related risks. Additionally, in August 2025, the DOJ announced a criminal resolution with a healthcare insurance company that misused AI by making improper payments to pharmacies for patient referrals submitted through the company's AI platform.[9]

Against this backdrop, it would be advantageous for companies to monitor the changing regulatory landscape and ensure a strong AI governance framework is in place. The DOJ's recent enforcement action as described above indicates that the future landscape of enforcement in this area is to be determined and, in the meantime, the DOJ's "Evaluation of Corporate Compliance Programs" reinforces the need for an effective AI governance framework that spans the entire lifecycle of the AI technology and is integrated into GxP, enterprise risk management and compliance frameworks, such as alignment with quality systems, pharmacovigilance processes and clinical trial oversight. Throughout this lifecycle, a variety of stakeholders across the company will need to be engaged to effectively evaluate and mitigate associated risks.

Management of Emerging Risks to ensure Compliance with Applicable Law (DOJ, ECCP (September 2024)):

- Does the company have a process for identifying and managing emerging internal and external risks that could potentially impact the company's ability to comply with the law, including risks related to the use of new technologies?
- How does the company assess the potential impact of new technologies, such as AI on its ability to comply with criminal laws?
- Is management of risks related to use of AI and other new technologies integrated into broader enterprise risk management (ERM) strategies?

- What is the company's approach to governance regarding the use of new technologies such as AI?
- How is the company curbing any potential negative or unintended consequences resulting from the use of technologies?
- How is the company mitigating the potential for deliberate or reckless misuse of technologies, including by company insiders?
- To the extent that the company uses AI and similar technologies in its business or as part of its compliance program, are controls in place to monitor and ensure its trustworthiness, reliability, and use in compliance with applicable law and the company's code of conduct?
- Do controls exist to ensure that the technology is used only for its intended purposes?
- What baseline of human decision-making is used to assess AI?
- How is accountability over use of AI monitored and enforced?
- How does the company train its employees on the use of emerging technologies such as AI?

AI Governance Framework

Life sciences companies should pull all AI-relevant strategic and risk management decisions together in a workable and user-friendly AI governance framework. The framework should align with the company's strategic mission, adapt to a variety of use cases, and help ensure the right people are involved at the right time to foster a balance between strategic vision, operational excellence, and risk management.

A comprehensive AI governance framework includes the activities necessary to manage the risks of a company's AI development and implementation. In addition to risk management, the framework should provide structure that helps ensure that the company expends resources on key activities or functionalities that support business goals.

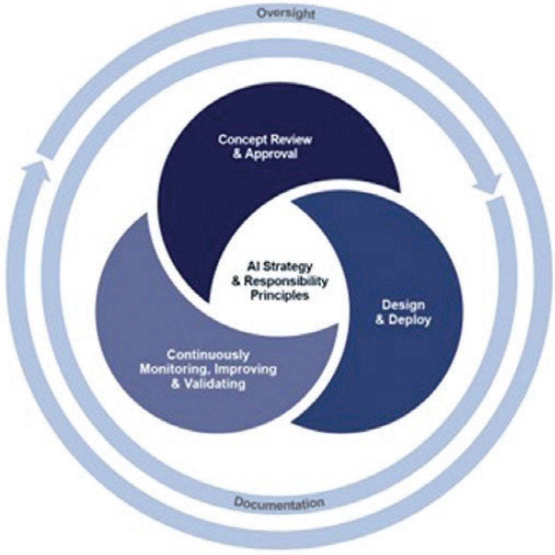
Oversight
Concept Review & Approval
AI Strategy & Responsibility Principles
Design & Deploy
Continuously Monitoring, Improving & Validating
Documentation

When developing and deploying an AI governance framework, life sciences companies should consider a three-stage approach:

- **Stage One: Concept Review and Approval**—Define how to bring together the right stakeholders to evaluate the balance between cost, benefit and risk of a given AI business use case, and set conditions for implementation that help ensure the desired balance will be realized in bringing the concept to reality. This stage can leverage concepts and learnings from the medical, legal and regulatory (MLR) review process, which brings cross-functional perspectives to assess risk and gather alignment before proceeding.
- **Stage Two: Design and Deploy**—Define the risk management and documentation standards for each AI model as it is developed, focused on regulatory expectations. Establish oversight to ensure it is developed as defined in Stage One and require reapproval for material changes. Existing quality and validation processes in life sciences share similarities with emerging AI governance needs. Companies can draw lessons from these established practices, particularly around documentation, traceability and

oversight to help meet requirements such as the EU AI Act's conformity assessments.

- **Stage Three: Continuously Monitoring, Improving and Validating**—Define how to establish a plan for business oversight and continuous testing for each AI model. Make sure the model is staying true to its intended business purpose. These ongoing monitoring requirements mirror practices in pharmacovigilance or post-marketing surveillance, where continued evaluation is required.

The decisions made at all stages must align with the company's AI strategy and principles of responsible use. The core of each stage includes robust documentation and oversight, which are essential elements that should be embedded throughout.

Setting the Tone: AI Strategy and Principles of Responsible Use

The OECD values-based principles provide recommendations for innovative and responsible use of AI within international organizations.

The OECD Values based principles address five key areas of responsibility for AI stakeholders:

1. Inclusive growth, sustainable development and well-being,
2. human rights and democratic values, including fairness and privacy,
3. transparency and explainability,
4. robustness, security and safety, and
5. accountability."[10]

The OCED principles have been widely accepted across the globe as the harmonized international governance approach. They have been adopted by the OECD's 70-member states, which includes the U.S. and the G20. The principles are the basis for the global regulations and standards including the EU AI Act and NIST's AI Risk Management Framework.

Organizations such as the Organisation for Economic Co-operation and Development (OECD), National Institute of Standards and Technology (NIST), and the International Organization for Standardization (ISO) have established standards for what constitutes responsible stewardship when developing, deploying, and using AI applications.[11] For example, the OECD standards include respecting the rule of law, committing to transparency and responsible disclosure, ensuring that AI systems are safe and secure against foreseeable misuse or adverse conditions, and maintaining traceability across datasets, processes and decision-making throughout the AI model's lifecycle. To meet these expectations, companies should consider internalizing the OECD principles and establishing their own principles around the responsible use of AI within their organization.

These principles should be right-sized and aligned with the purpose and business objectives of the company. When defining these principles, companies should consider factors such as the company's:

- Size and complexity
- International presence
- Regulatory exposure
- Mission and values
- Business/operating model
- Competitive positioning

Embedding Accountability: Business-Led Executive Sponsorship

Before an AI use case enters the concept review and approval stage, it must be vetted by the initiating business function. This includes securing alignment among functional leadership and designating an individual executive sponsor who will maintain oversight throughout the governance lifecycle.

The executive sponsor will be responsible for:

- Ensuring the AI use case aligns with the company's overall business strategy
- Prioritizing resource allocation based on trade-offs between risk and business benefit

- Maintaining accountability for the AI model and ensuring it remains aligned with the initially defined purpose

As a consequence of operating in a complex regulatory environment, most life sciences companies already have a strong compliance foundation to leverage with robust policies and processes against the backdrop of strong, established cultures of compliance and ethical decision-making as part of their mission and values. These existing frameworks and mindsets will help guide executive sponsors and functional leadership in adopting and aligning to established AI principles for responsible use.

Building on this foundation, effective AI governance requires oversight that extends beyond the initial approval and deployment of AI use cases. This is especially true in the context of AI applications, given the inherent nature of its biggest benefit—the ongoing adaptation and refinement of the AI model. As a result, oversight of the model will look different at each stage in the lifecycle; however, the executive sponsor will be accountable throughout the lifecycle.

Establishing Trust and Defensibility: Robust Documentation

"Effective risk management and oversight of AI hinge on a critical, yet underappreciated tool: comprehensive documentation."[12] Documentation does more than satisfy regulator expectations; it enables organizations to identify weaknesses across the AI model's lifecycle and supports real-time course correction.

Each business need and associated use case must clearly define the AI model's purpose, what information/data will be used, and the desired output. These elements and the decisions surrounding them—including a strong intent for the use of AI—must be formally documented in a Purpose and Request Form. To initiate stage one (concept review and approval) of the AI governance framework, a business need with strong intent for the use of AI must be documented and signed off by the executive sponsor. This Purpose and Request Form is not only important for the initial concept review and approval, but is also referred to in all stages of the lifecycle. Similar to the Third-Party Needs Assessment Form, the Purpose and Request Form will provide the necessary documentation of the business justification and required approvals.

Potential Elements of the Purpose and Request Form:

- Purpose and Strategic Alignment
- Scope (Global, etc.)
- Input Data
- Desired Output
- Potential Tools (if external)
- Cost & Budget Implications (both development and maintenance)
- Business Benefits vs. Risk Trade-offs
- Monitoring & Oversight Plan

"In essence, documentation is a tool that has the potential to—and is indeed necessary to—facilitate both external accountability and internal risk management practices."[13] Once a model is approved and advances into the design and deploy stage, documentation plays a key role in providing transparency: explaining how the AI model arrives at certain decisions, identifying the data used to train the algorithms, reconfirming the intended use cases, and supporting the process of risk management and mitigation.

As the AI model progresses through its lifecycle, documentation also supports oversight functions in confirming that the model continues to operate as intended and remains aligned with the originally defined purpose.

What's Next? Stage One: Concept Review and Approval

Once a Purpose and Request Form has been completed and receives both functional approval and executive sponsorship, a series of gatekeeping activities should occur before the AI model can be deployed. We refer to these gatekeeping activities as ***Stage One:***

Concept Review and Approval. At this stage, various cross-functional committees should evaluate the rationale and proposed AI use case across multiple perspectives, including strategic and business alignment, financial and technical feasibility, and legal and reputational risk.

This structured yet flexible AI governance framework can be adapted to different AI use cases across various functions within the organization, each carrying varying levels of risk. Knowing how to calibrate the governance process for each of these variables will help place the right subject matter experts at the right place at the right time to weigh in and help design a model that is aligned with the intended purpose.

To support this calibrated approach, life sciences companies can draw on their experience navigating a complex regulatory environment, including opportunities to leverage existing governing bodies, tools, and processes when implementing and defining an AI Governance Framework.

In our next article of this series, we will describe how adapting the conceptual approach of processes such as the MLR review can help build and operationalize AI governance.

Endnotes

1. "Generative AI is a subfield of AI in which computer algorithms are used to generate outputs that resemble human-created content, such as images, videos, art, music, text, and software code. The output is based on training data inputted into large models—often containing millions of images, sentences, and/or sounds—from which a computer can learn to create the desired output." Avi Weitzman & Jackson Herndon, Generative AI: The Next Frontier for Section 230 of the Communications Decency Act, N.Y.L.J. (June 26, 2023) available at *https://assets.ctfassets.net/t0ydv1wnf2mi/4LqDaZLG1NvNeWhAoGLxRq/f0f003214772a4c67369de054a9870b8/NYLJ706202344999Paul.pdf* (last accessed May 14, 2025).
2. 15 U.S.C. § 9401(3), available at *https://www.govinfo.gov/content/pkg/USCODE-2020-title15/html/USCODE-2020-title15-chap119.htm* (last accessed Sept. 17, 2025).
3. Dan Huttenlocher, Asu Ozdaglar & David Goldston, A Framework for U.S. AI Governance: Creating a Safe and Thriving AI Sector (MIT Schwarzman Coll. of Computing, Nov. 28, 2023), available at *https://computing.mit.edu/wp-content/uploads/2023/11/AIPolicyBrief.pdf* (last accessed Oct. 10, 2025).
4. Camille Paulhac & Jason Raeburn, European Commission & AI: Guidelines on Prohibited Practices (Apr. 3, 2025), Paul Hastings, available at *https://www.paulhastings.com/insights/client-alerts/european-commission-and-ai-guidelines-on-prohibited-practices* (last accessed Sept. 5, 2025).
5. EXEC. OFF. OF THE PRESIDENT, America's AI Action Plan (July 23, 2025), available at *https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/* (last accessed Sept. 5, 2025).
6. On September 10, 2025, Senator Ted Cruz proposed the SANDBOX Act (S. 2750), which would establish a process for AI deployers or developers to apply for temporary modifications or waivers of federal regulations. *See* U.S. Senate Committee on Commerce, Science, & Transportation, Sen. Cruz Unveils AI Policy Framework to Strengthen American AI Leadership (Sept. 10, 2025), available at *https://www.commerce.senate.gov/2025/9/sen-cruz-unveils-ai-policy-framework-to-strengthen-american-ai-leadership* (last accessed Oct. 15, 2025).
7. *See, e.g.*, U.S. FOOD & DRUG ADMIN., Draft Guidance: Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations (Jan. 7, 2025), available at *https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing* (last accessed Sept. 5, 2025); U.S. FOOD & DRUG ADMIN., Final Guidance: Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Function (originally issued Dec. 4, 2024; issued on Aug. 18, 2025), available at *https://www.fda.gov/regulatory-information/search-fda-guidance-documents/marketing-submission-recommendations-predetermined-change-control-plan-artificial-intelligence* (last accessed Sept. 5, 2025).
8. 45 C.F.R. Parts 170 and 171, HTI-1 Final Rule, Effective Feb. 8, 2024, available at *https://public-inspection.federalregister.gov/2024-29163.pdf* (last accessed Sept. 3, 2025).
9. Press Release, U.S. Dep't of Justice, National Health Care Fraud Takedown Results in 324 Defendants Charged in Connection with Over $14.6 Billion in Alleged Fraud (June 30, 2025), available at *https://www.justice.gov/archives/opa/speech/deputy-attorney-general-lisa-o-monaco-delivers-remarks-university-oxford-promise-and* (last accessed Sept. 18, 2025).
10. Org. for Econ. Co-operation & Dev., "AI Principles," OECD (n.d.), *available at https://www.oecd.org/en/topics/ai-principles.html* (last accessed Oct. 15, 2025).

11. ORG. FOR ECON. CO-OPERATION & DEV., Recommendation of the Council on Artificial Intelligence, OCED/LEGAL/0449 (Mar. 5, 2024), available at *https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449*; National Institute of Standards and Technology, AI Risk Management Framework (Jan. 26, 2023), available at *https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf*; International Organization for Standardization, ISO/IEC 42001:2023 (Dec. 2023), available at *https://www.iso.org/standard/42001* (last accessed Oct. 15, 2025).
12. Amy Winecoff & Miranda Bogen, Best Practices in AI Documentation: The Imperative of Evidence from Practice, CTR. FOR DEMOCRACY & TECH (July 25, 2024), available at *https://cdt.org/insights/best-practices-in-ai-documentation-the-imperative-of-evidence-from-practice/* (last accessed Oct. 15, 2025).
13. *Id*.

Reprinted from Journal of Health Care Compliance, Volume 28, Number 2, March–April 2026,
pages 17–22, 60–61, with permission from CCH and Wolters Kluwer.
For permission to reprint, e-mail permissions@cch.com.