

Arnold & Porter

The Convergence of Life Sciences and Artificial Intelligence

Seizing Opportunities While Managing Risk

NOVEMBER 2024



Contents

Section I: Survey

- 3 Executive Summary**
- 5 Key Findings
- 6 Developing AI Use Cases in Life Sciences**
- 7 AI and the Life of a Biomedical Product**
- 8 AI's Potential Across the Life of a Biomedical Product
- 9 Product Discovery and Design
- 12 Product Development
- 16 Product Manufacturing, Commercialization, and Stewardship
- 21 AI Governance and Compliance**
- 24 Using AI in Regulatory and Compliance Functions
- 26 Survey Conclusion**

Section II: Resources

- 27 Understanding and Managing the Life Sciences/AI Convergence**
- 29 The Developing Framework for Global AI Regulation**
- 32 Key Regulatory Considerations in Adopting AI Technologies**
- 36 Questions to Ask Before Deploying AI
- 37 AI Privacy and Cybersecurity Risks**
- 38 IP Risks in AI and Life Sciences**
- 40 Promotion of AI Use and Avoiding "AI Washing"**
- 41 AI in Litigation**
- 43 Getting Paid for AI-Related Services and Tools**
- 46 Mitigating Risk in Negotiating AI Vendor Contracts

- 48 Glossary**
- 49 Methodology**
- 49 About Arnold & Porter**
- 50 Acknowledgements**

Survey and Resources

As detailed in this report, AI technologies present significant legal risks and issues. This report should not be construed as providing legal advice or a legal opinion on any specific facts or circumstances in any jurisdiction. The contents are intended for general informational purposes only, and you are urged to consult with counsel on any specific legal questions you may have.

Executive Summary

“Every new technology is a double-edged sword, but AI may be the sharpest blade yet.”

– Deputy Attorney General Lisa Monaco, speaking in February 2024.

The life sciences industry is in many ways converging with the field of artificial intelligence (AI). Whether companies are developing therapeutics, vaccines, medical devices, or diagnostics, life sciences companies are actively exploring the incorporation of the technology into products, as well as critical operations throughout the product lifecycle, from discovery through commercialization.

This convergence is highly complex, and made even more so by the volatile state of the AI legal, policy and political debate around the world. Indeed, as we are finalizing this report, Donald Trump has been re-elected and it is likely that the executive order on AI issued by President Biden will be withdrawn next year, reportedly to be replaced with a more innovation-friendly approach. As the industry knows from the past, however, governmental enforcement restraint does not necessarily follow such policy changes.

Much of current AI activity in the life sciences industry is focused on discovery and research efficiency as well as optimization of operations to reduce time and costs. However, a major driver over time will be navigating the role of AI in the healthcare system of the near future, when AI tools and algorithms will likely drive a significant portion of patient care.

Increasingly the patient journey will also be an AI journey — they will be diagnosed using imaging or digital tools developed with AI, treated using a sophisticated AI algorithm, enrolled in an AI-designed and monitored study, transitioned to the approved product utilizing an AI-customized treatment plan, and then be subject to monitoring via wearable sensors or implants incorporating AI. Each such healthcare development will have profound implications for those — including the life sciences industry — dedicated to serving these patients.

Yet, as noted, these developments are not without risk. As life sciences companies evaluate and adopt predictive AI, generative AI, and machine learning (ML) models, these new technologies also open companies to a rapidly evolving AI regulatory and enforcement landscape that poses compliance challenges for the much-scrutinized life sciences industry.

To understand how life sciences companies are addressing this emerging and disruptive technological convergence, Arnold & Porter surveyed 100 senior industry executives and department heads in key technology, leadership, and compliance roles. Respondents represented biopharmaceutical, digital health, medical device, and diagnostic companies, among other types,

based in North America and Europe; the majority reported between US\$101 million and US\$50 billion in 2023 gross revenue.



Our research found that while the adoption of AI is relatively new for many life sciences companies — three-quarters of our respondents only began implementation in the past two years — the pace

of adoption is rapidly accelerating. **Eighty-six percent of organizations are now in the process of implementing plans to deploy AI use cases within two years or less for research and development (R&D), manufacturing, marketing, regulatory, and other applications.**

However, as is often the case with transformative technological change, governance measures appear to be lagging behind AI implementation. While many companies may be waiting to get more experience with piloting AI use cases, just 55% of respondents that are currently using AI have put AI policies and standard operating procedures in place.



That process is likely to accelerate greatly in the coming year. While incredibly promising for patients and industry, the use of AI in life sciences can heighten compliance risks given the close link between some AI tools and clinical decisions, patient safety, and ethical treatment, as well as the use of huge volumes of often sensitive data. The cost and complexity of AI can also create major challenges by reducing transparency, accountability, and equity, and may heighten fraud and abuse risks in the deployment of AI tools and strategies with healthcare practitioners and patients.

Thus, in parallel with the patient journey shaped by AI, the industry will need to continually assess and control its activities to ensure that the risks from the incorporation of AI into the business do not produce investigations or litigation, delay important transactions, or embed financial vulnerabilities that could severely damage the enterprise in the years to come.

In what follows, we'll examine key industry benchmarks for AI implementation across the product lifecycle, including detailed insights on how companies are progressing in AI adoption, governance, compliance, and risk mitigation strategies. In the last section of this report, we also provide Arnold & Porter resources on some of the key risk considerations that companies must address in parallel with this fundamental technological and healthcare evolution.

A glossary is provided on page 48.

Key Findings

75%



AI implementation is still in the early stages for many life sciences companies — but adoption is accelerating. Three-quarters (75%) of respondents began implementing AI within the last two years; only 5% started more than five years ago. However, 86% of those now in the process of implementing AI tools are planning to deploy them in the next two years or less.

79%



Life sciences companies are adopting AI across the product lifecycle. Almost eight in 10 (79%) life sciences companies surveyed are using or planning to use AI in R&D, with significant segments planning to or already implementing AI for manufacturing (62%), marketing (45%), regulatory (42%), and compliance (29%) functions.

74%



Intellectual property (IP) presents significant concerns for life sciences companies using AI. Nearly three-quarters (74%) of respondents have a high level of concern about facing AI-related IP issues in the upcoming year; only 3% express no concern at all.

51%



AI is transforming product discovery and design. Approximately half (51%) of respondents have explored leveraging AI to optimize the product discovery and design process, citing improved efficiency and faster time-to-market for new products as key benefits. Sixty-three percent are using or plan to use AI-driven testing and simulations.

85%



AI is optimizing marketing and sales strategies. The majority (85%) of respondents report that AI-driven initiatives to optimize sales strategies are highly effective in the product commercialization and stewardship phase.

55%



Governance appears to be lagging behind AI implementation. Just 55% of respondents currently using AI have put AI policies and standard operating procedures in place; even fewer have completed regular AI audits or assembled cross-functional teams to promote safe and effective use.

Developing AI Use Cases in Life Sciences

Throughout this report we reference AI “use cases” and “tools” of various types. Many AI uses are quite familiar, such as certain clinical decision support tools and chatbots. However, the capabilities of these tools are growing rapidly, and use cases incorporating generative AI present particularly powerful applications that require risk assessment. Some examples include:

Retrieval augmented generation (RAG) is an AI framework that combines **extractive AI** (which retrieves reliable information from trusted sources) with **generative AI** (which refines or generates the text based on the trusted information). RAG models are gaining traction for their ability to improve the **accuracy** and **attributability** of AI-generated content by grounding responses in specific, factual data.

RAG can be especially useful in life sciences, where accuracy and transparency are paramount. For example, it can help in **drug discovery and development**, where pulling from validated external datasets (e.g., clinical trial data, peer-reviewed studies) before generating insights ensures that AI-driven conclusions are grounded in reliable information. RAG also helps in regulatory compliance, where ensuring traceability is key.

Generative AI combined with human-assisted machine learning (HAML) is another breakthrough that is enabling more sophisticated analysis of both **structured and unstructured data**. This combination allows AI systems to combine insights from unstructured sources (like medical images or physician notes) with structured data (like clinical trial results or patient demographics), creating a more comprehensive view of the data.

In fields such as **precision medicine** and **clinical trials**, this combination allows for better uses of different kinds of data. With the combination of generative AI and HAML, AI in life sciences is evolving beyond simple predictive models to incorporate **multimodal insights** that improve patient outcomes, speed up R&D processes, and drive innovations in **personalized healthcare**.

Agentic workflows refer to the use of AI agents that autonomously perform complex, multistep tasks, adapting in real time based on new data inputs. These AI systems could be used in the future to carry out tasks like clinical trial management, drug manufacturing, or patient monitoring with minimal human intervention, making them ideal for high-stakes, data-driven environments.

In drug manufacturing or clinical trials, agentic workflows can manage logistics, adjust parameters in real time, and make data-driven decisions that reduce human error and improve efficiency. They can also integrate data from multiple sources to inform actions, allowing for a **dynamic, adaptive approach** to complex processes.

AI and the Life of a Biomedical Product

While AI tools have the potential to transform an array of functions within the product lifecycle, most life sciences companies are still in the early stages of evaluation and implementation. Just 5% of respondents we surveyed had started to implement AI more than five years ago, with the majority (75%) only beginning that process in the last two years. Yet the industry appears poised to accelerate its AI usage in the near term, as 86% of companies currently integrating AI expect to deploy it in product lifecycle functions in two years or less.

Life sciences companies are exploring AI for a wide range of functions across every stage of the product lifecycle, from optimizing trial design and recruitment to strengthening marketing and sales strategies. Improving the product design and discovery process is the most typical initial focus, with slightly more than half (51%) of respondents exploring opportunities in this area; respondents also see promise in leveraging AI for data privacy (48%) and analytics functions (42%).

Respondents almost uniformly predict substantial benefits from the industry-wide adoption of AI over the next five years. Nine in 10 respondents agree or strongly agree that widespread AI adoption will make products more cost-effective for patients and the healthcare system, while 86% say the same of AI's ability to help companies get products to market — and thus to patients —

faster. These efficiencies extend internally too: 85% agree or strongly agree that industry-wide AI use will yield cost efficiencies for companies. Additionally, and perhaps surprisingly, more than two-thirds (67%) of respondents say that the life sciences industry's use of AI will actually *deepen* trust among consumers.



In the following sections, we will explore in-depth the different use cases, opportunities, and concerns that AI tools present for the life sciences industry, at every stage of the product lifecycle.

AI's Potential Across the Life of a Biomedical Product

Increasingly, AI-driven tools will be pervasive in delivering diagnoses and treatments to patients.



Diagnostic, imaging, and digital tools developed with or utilizing AI will assist (and in some cases largely replace) healthcare practitioners in detecting and identifying patient conditions; sophisticated AI algorithms will then recommend next steps to either treat the patient or further confirm the nature of the disease.



Patients with conditions where no adequate therapy exists will enroll in clinical trials under AI-designed protocols involving therapies developed with AI tools and tested in part via virtual tests and patient digital twins powered by AI.



AI tools will also assist in designing clinical trials, analyzing trial data, and compiling regulatory filings. Digital biomarkers will be used to assess patients during trials, monitored via AI tools.



Product manufacturing and supply and distribution chains will also be AI-driven, resulting in more efficient global production, allocation, and tracking that can reduce the risks of shortages.



Upon product approval, patients will transition to commercially available products, utilizing AI-customized treatment plans, with reimbursement and patient assistance also AI-assisted.



AI will assist in patient care via monitoring systems and wearable sensors or implants that track their condition and medication adherence. Healthcare practitioner alerts will allow for rapid responses to changing patient conditions.



Manufacturers will utilize AI to find additional healthcare practitioners likely to have patients with approved conditions, then use AI tools to educate them further on diseases and treatments as well as market products via AI-optimized advertising and social media efforts.

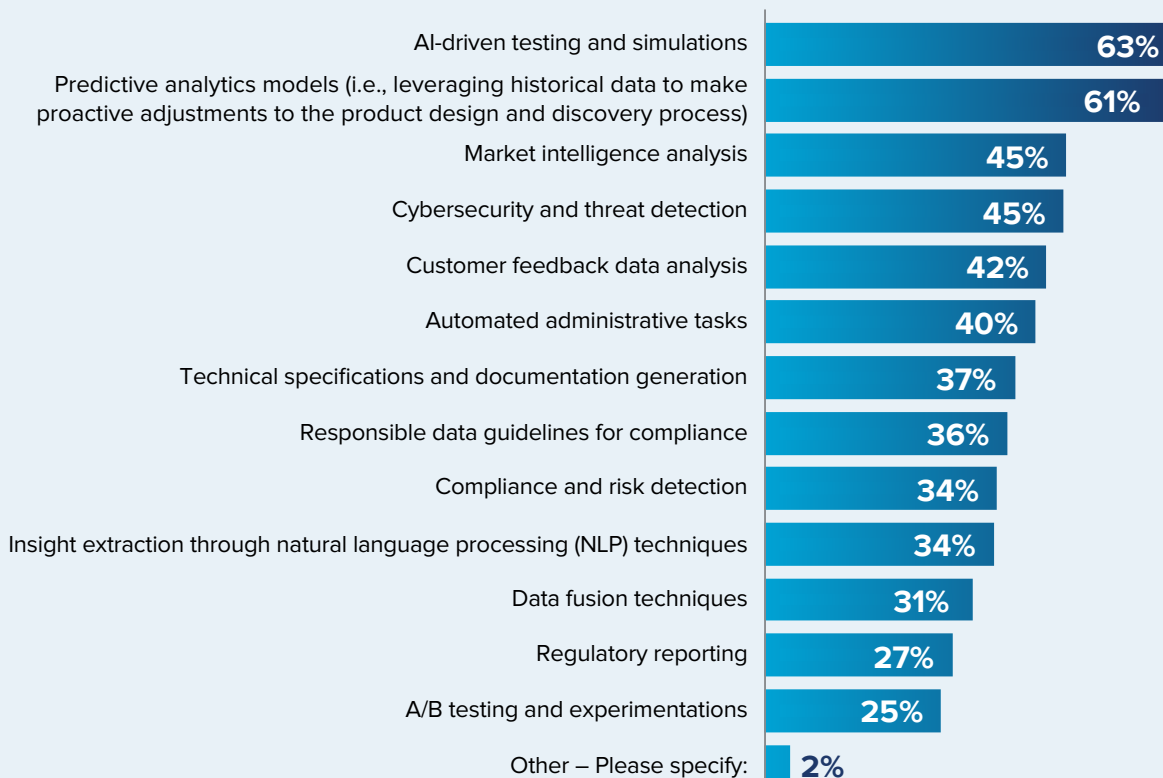
Product Discovery and Design


Of all the stages of the life of a product, the time-consuming, expensive, and often frustrating process of creating, identifying, or designing new drugs and devices is currently among the most promising and immediate areas of focus for AI applications. Drawn by the prospect of faster timelines and cost savings, life sciences companies are making significant investments in this area. Nearly eight in 10 (79%) respondents have either already implemented or are planning to implement AI in their R&D functions.

Tools like generative AI could have a transformative impact on the discovery of both small molecules and biologics by synthesizing data about molecules and protein structures to generate new therapeutic product candidates, potentially shaving years off the often decade-plus, multibillion-dollar process of bringing a new drug to market. Coupled with advances in automation, some AI tools can rapidly probe real-world data to improve drug candidates' designs, while others focus on computationally predicting interactions to reduce the need for screening thousands of different molecules in the lab.

Indeed, for the life sciences companies we surveyed, AI-driven testing and simulations are the most popular product design and discovery use cases, selected by 63% of those implementing AI in R&D. Predictive analytics models, such as those leveraging historical datasets or clinical trial information, are close behind, at 61%. Such use cases have enormous implications for avoiding unnecessary expenditures of resources and time in the always risky pathway to developing a new treatment.


Potential use cases of AI in product design and the discovery phase



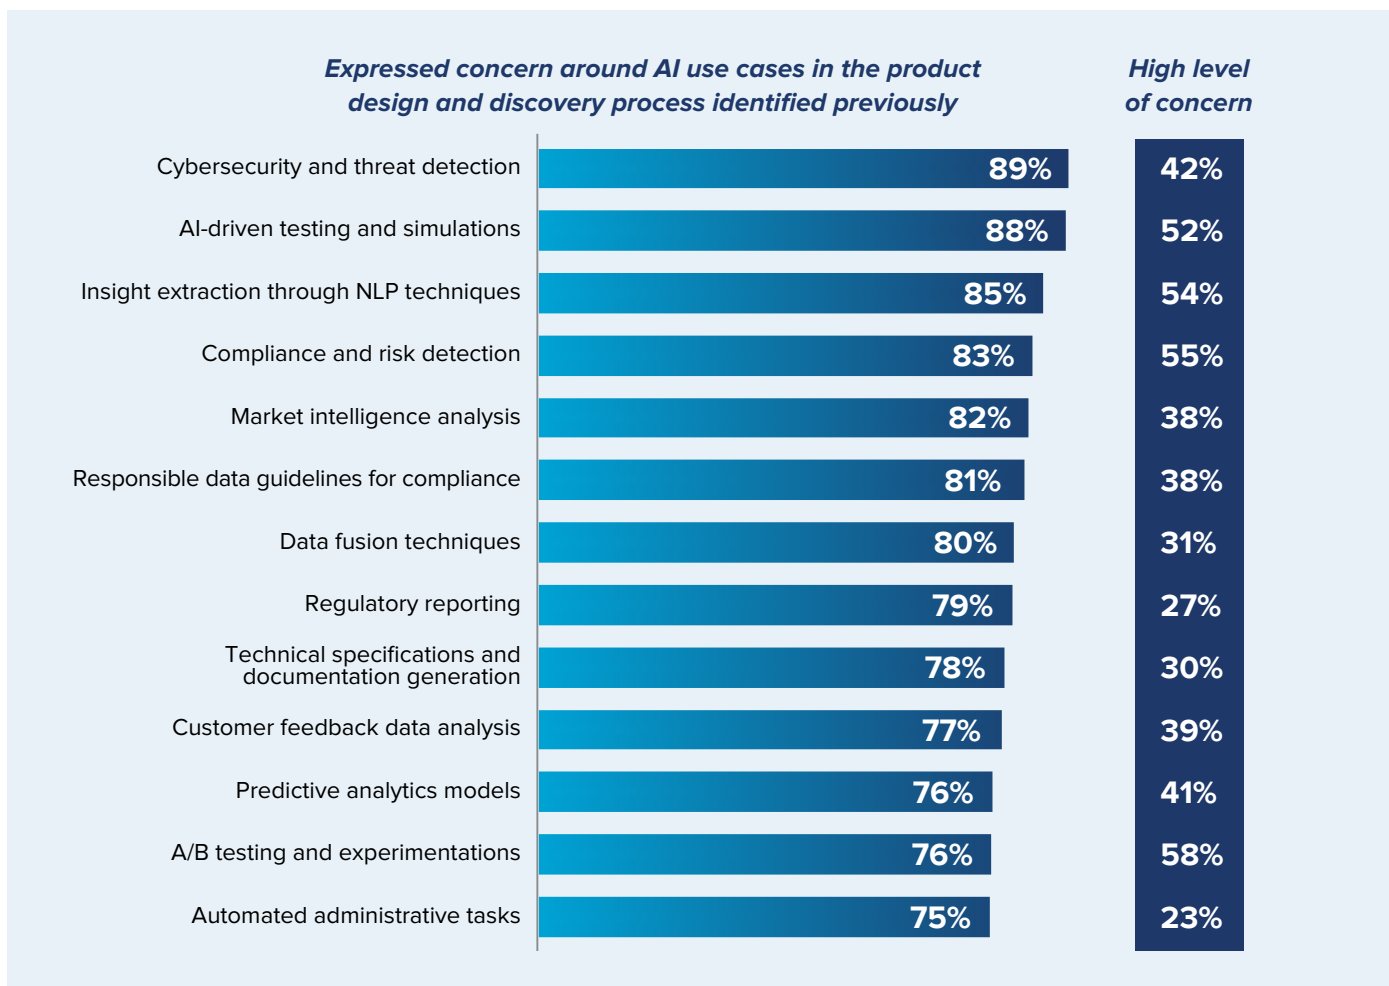


Efficiency is a key objective for AI in product design and discovery. More than half of respondents point to improved efficiencies in the product development process (61%) and a faster time-to-market for new products (57%) as the most significant advantages that AI usage offers during this stage. Yet improvements to product design (56%) and enhanced quality control (52%) are also among the top reasons cited.

Opportunities that will arise from AI usage in product design and the development phase



Yet, as with any groundbreaking technology, life sciences companies have hesitations about using AI for these important functions. When it comes to factors related to using AI for product design and discovery, half of respondents are highly concerned about liability issues, and even more (58%) are significantly concerned about privacy and confidentiality. The vast majority (88%) of those using or planning to use AI for testing and simulations also expressed concern about leveraging AI, with 52% saying they have a high level of concern. In fact, there are concerns about use cases across the board: every AI application was a worry for at least 75% of the respondents who were using or exploring them. Much of this anxiety over AI is driven by the unknowns in the application and regulation of AI, as well as uncertainties about best practices for avoiding negative outcomes.



While the potential benefits of AI to speed and enhance product design and discovery are substantial, life sciences companies appear to be wrestling with the formidable challenges of bringing emerging technologies like generative AI into such a foundational aspect of their business — innovation.

Product Development


AI also holds promise for life sciences companies to gain an edge in the preclinical and clinical development stages. Nearly two-thirds of respondents surveyed expect AI integration will lead to long-term cost savings (63%) and accelerated product development (61%). Significant segments also anticipate AI will be able to automate administrative tasks (44%) and accelerate clinical trials (43%).

Advantages that AI integration will bring to product development processes



For example, AI tools could assist in clinical trial protocol design by enabling a sponsor to hone in on the optimal sample size, dosing strategy, and patient population. AI can also assist in identifying and validating study biomarkers and clinical endpoints, which are critical to the acceleration of innovation. Currently, 47% of respondents are using or exploring the use of AI for virtual patient models.

Since patient recruitment can be one of the slower and more resource-intensive elements of clinical research, many major pharmaceutical companies have already adopted AI to identify potential trial subjects — and clinical investigators. Our survey found that 44% of respondents are utilizing AI for subject onboarding and retention. Once the trial is underway, AI can also support safety monitoring and optimize the use of real-time monitoring tools such as wearable devices. In fact, among the life sciences companies we surveyed, AI-enhanced diagnostic and monitoring tools are the second most popular application currently in use or under consideration at this stage in the product lifecycle, at 59%.

Current or planned utilization of AI to improve the product development process

Another valuable use case is leveraging AI to manage, sort, clean, and probe large volumes of data and images; this includes data across multiple studies and sources, such as unstructured real-world data. Whether drawn from registries or electronic health records, such real-world evidence comes with myriad risks, from ensuring the quality of data to addressing compliance issues such as privacy and the fair market value of transfers to healthcare institutions and other data holders.


Notably, however, just a quarter of respondents say they are using or intending to use AI for pharmacovigilance or other safety reporting needs, indicating that companies are proceeding carefully in this area. In fact, 61% of those same respondents who are using or planning to use AI for pharmacovigilance are highly concerned about such a use case — more so than for any other answer choice — suggesting the risks are very high when it comes to this critical safety function.

Companies appear to be cautious in utilizing AI in clinical trial-related compliance and risk detection activities, likely moving carefully given established models for oversight and extensive procedures already in place in this area. Given the risks involved in using AI in research, it will be critical for sponsors to develop AI-specific internal policies, for quality control and monitoring plans, among other things, and ensure adequate training and resources to validate AI applications to ensure they are fit for purpose and reproducible.



- **For example, clinical trial sponsors are responsible for ensuring bias is controlled in trial design, data collection, and results analysis.** An AI/ML tool that has been trained on a sponsor-selected dataset may be unable to properly discern patterns that are inconsistent with the AI's trained experience, leading to an output unduly favoring a particular outcome.
- **AI tools may also utilize sensitive patient data, requiring careful attention to data security and patient confidentiality.** Large amounts of training data may also be required, which could violate patient privacy or create security risks.
- **AI use in clinical trials must also be inspectable by regulators, requiring access to vendor validation.** Companies must ensure appropriate support from AI vendors in addressing areas of regulatory scrutiny with respect to AI-based submissions and decision-making.
- **Companies using AI for clinical trials to manage patient data or analyses also must consider avoidance of infringement of IP rights and compliance with data protection rules,** as well as with new AI-specific legislation such as the EU AI Act, which entered into force on August 1, 2024, and which impacts any company offering goods or services using AI anywhere in the supply chain in the EU.

If sponsors intend to successfully implement AI applications in clinical development programs, understanding and managing risks through robust governance, quality control, and data management practices — from protocol design to final submissions — will be critical. Regulators are just beginning to grapple with the challenges of regulating AI in clinical research, and stakeholders should carefully monitor and help shape regulator policy-making in this area to both ensure workable standards and the practicability of implementation so the full benefit of AI in research can be achieved for patients and the healthcare system.



Product Manufacturing, Commercialization, and Stewardship

Assuming the AI-assisted development of a product is successful, AI will also be critical in the process of bringing the product to patients at scale, from the distribution chain to patient support to safety and compliance monitoring. AI will also be an important tool in product commercialization: nearly half (49%) of respondents are already using generative AI in marketing functions, for example, and 57% are in the process of implementing predictive AI tools to help sell products. The most popular use cases being utilized or in development in the product commercialization phase include market intelligence analysis (57%) and product data analysis (56%), together with sales forecasting (55%).

How organizations are using or planning to use AI in the product commercialization phase

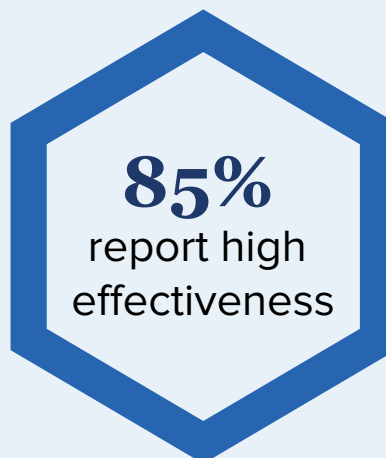


More than half (56%) of respondents see precision marketing as a key advantage of integrating AI into the commercialization process; for example, manufacturers could analyze data to locate healthcare practitioners with patients likely to have certain conditions and deploy AI tools to educate on diseases and treatments. Other advantages include enhanced customer engagement (49%), product and price optimization (48%), and product differentiation (47%). Overall, these tools appear to be having an impact — 85% of respondents say their organizations are highly effective at leveraging AI-driven initiatives to maximize sales.


Advantages of integrating AI into product commercialization processes



On a scale of 1 to 10, how effectively is your organization leveraging AI-driven initiatives (e.g., consumer support, product distribution) to maximize sales?



Life sciences companies are rightly cautious when considering AI for higher-risk applications at this stage. Every single respondent using or planning to use AI in incentive compensation, for example, reports concerns about these use cases. The same is true for those using or considering AI to manage product liability claims — with more than half (53%) *highly* concerned. Working with technology, legal, and compliance functions to navigate these risks and implement safeguards is an important part of responsible AI usage, particularly in the highly scrutinized life sciences industry.




As companies develop products and move toward global manufacturing, they are also turning to AI tools to improve product manufacturing and supply chain management. Demonstrating the technology's transformative potential, the majority (62%) of respondents say they have already implemented or are planning to implement AI functions in manufacturing; just 2% say they are not considering the use of generative AI or ML models for these ends.

The top two anticipated manufacturing benefits of AI use are increased productivity (60%) and streamlined procurement and logistics processes (58%). Approximately half of respondents say that AI will also enhance inventory management (51%) and increase quality assurance (49%), allowing for better monitoring and management of finished product inventories and shipments globally.

Quality control is overwhelmingly the most popular AI application in manufacturing, with two-thirds (66%) of respondents using or planning to use it that way. It is also the top use case for supply chain management, with 69% of respondents using or exploring AI for quality control functions — significantly more than any other supply chain application.

Quality control applications include using AI tools to inspect individual products for defects — such as cracks or discoloration issues that might be invisible to human inspectors — by analyzing images or video. Such tools can also check for inaccurate labeling or monitor environmental conditions during the manufacturing process to proactively anticipate defects. AI can perform these functions much faster and often more reliably than human inspectors, boosting accuracy and productivity.




Other major AI manufacturing applications include optimizing product processing parameters, demand forecasting, raw material supply chain management, modeling manufacturing process changes, regulatory “track and trace” requirements, monitoring complaints and deviations, and facilitating recalls and corrective actions, among others.

Yet, as with any analytical tool, useful outputs depend on quality inputs — a concern for a significant segment of respondents managing data related to post-market complaint processing and adverse event reporting. Given the staggering amounts of data used and generated by these AI applications, data quality, accuracy issues, and data integrity are all top challenges, selected by 43%, 31%, and 27% of respondents, respectively.

In terms of managing data, regulatory compliance (29%) and validation for regulatory purposes (25%) are also key issues respondents are facing or anticipating in the near term — which tracks, given the close regulatory oversight of manufacturing in the life sciences industry. Regulators will expect

companies to meet data quality and integrity standards when using AI, despite uncertainty surrounding how life sciences companies should apply existing manufacturing best practices and standards to these evolving technologies. Regulators may also want to inspect AI for validation and verification purposes, which could prove challenging. This is a common issue for life sciences companies using third-party vendors and cloud data, and will require careful attention to contractual provisions on supporting and demonstrating regulatory compliance.

In response to these challenges, life sciences companies are turning to a variety of risk mitigation measures, with no one option a clear winner, suggesting this area is very much still a work in progress. About half of those already using AI are vetting vendor AI procedures (54%), proactively updating all master agreements (50%), conducting regular AI system monitoring and evaluation (50%), and vetting vendor security measures (50%). However, just under half (47%) have put in place data encryption and access controls.




AI Governance and Compliance

For all the potential advantages AI holds, it also requires careful management and responsible usage. Yet, likely reflecting the evolving compliance environment in this area, there is a lack of consensus among life sciences companies we surveyed when it comes to ensuring compliance with diverse regulatory and compliance requirements. No single compliance strategy is in use by more than half of respondents, with just 46% saying their organizations are developing AI governance or management frameworks. Just 38% of respondents collaborate with legal experts for AI compliance functions, and only 37% monitor AI in their compliance programs, reflecting the uncertainty surrounding AI compliance best practices.

Among companies that have already implemented AI, just over half (53%) have AI policies and standard operating procedures in place. And, as the use of AI becomes ubiquitous, organizations in the field may not be fully aware of certain ongoing risks, as even fewer (51%) have completed regular audits and assessments within their organizations to understand which AI tools are in use and for what purposes.

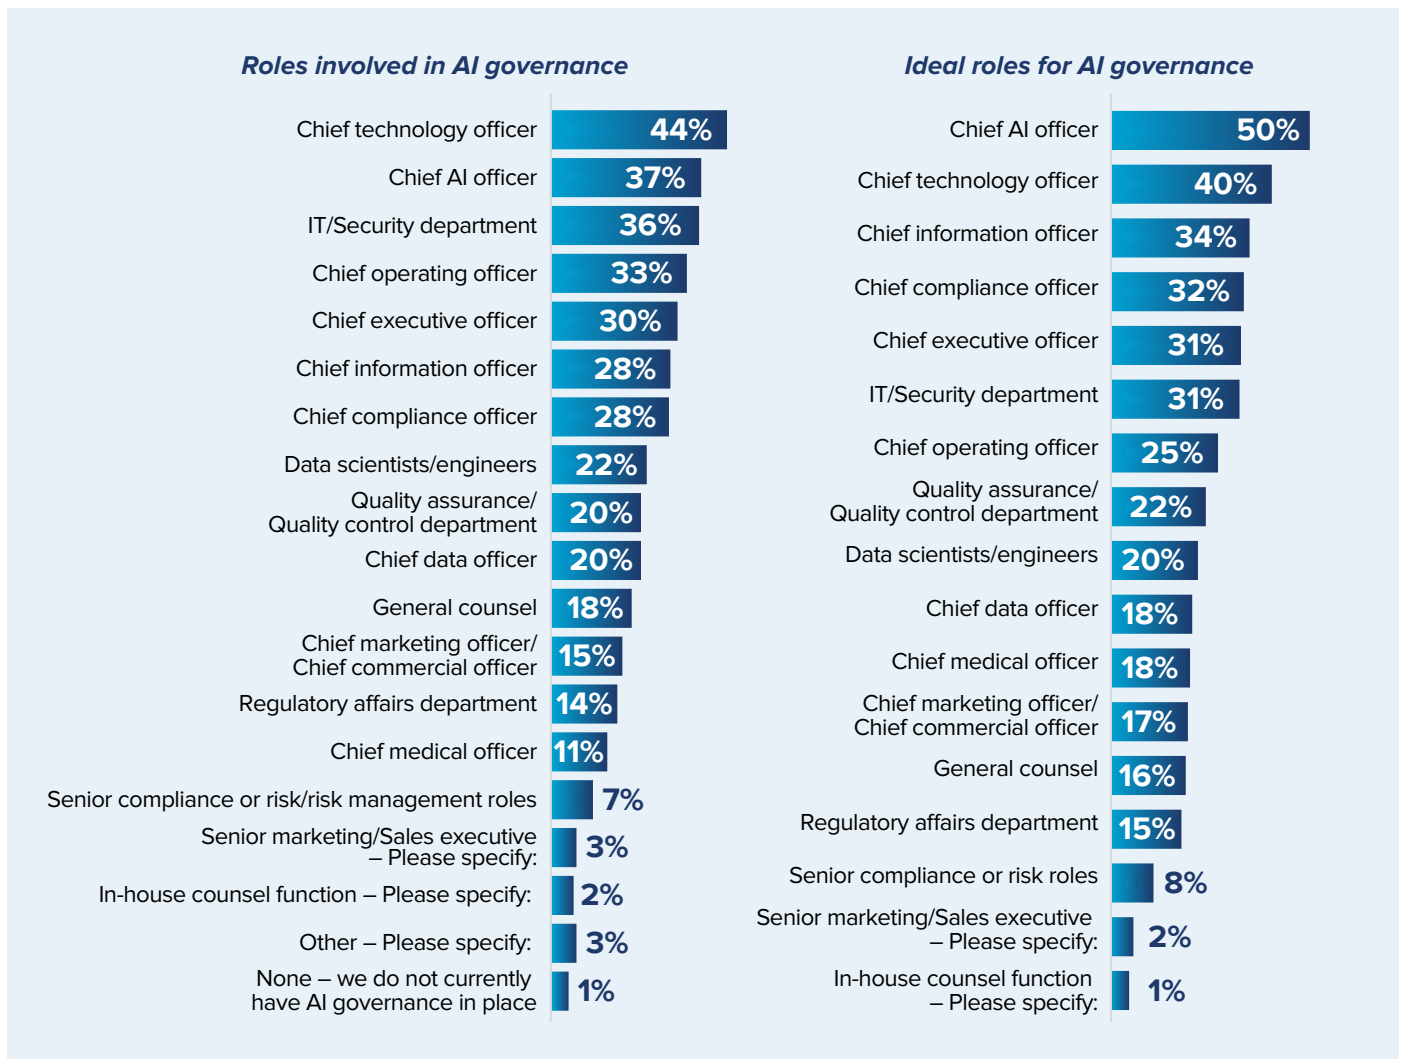
*Policies that organizations currently have in place as part of an AI governance plan
“Already in place/Using now” responses only*



Even the most straightforward governance steps, such as routine board discussions and agendas covering AI or enhancing transparency and accountability measures, are not universally in place among AI adopters. Fewer than half of AI-using respondents say they have assembled cross-functional teams or task forces, which can help leadership understand the risks and can help put these important policies in place.

In fact, just 44% of respondents say that chief technology officers are involved in AI governance within their organizations, potentially excluding an important viewpoint with deep knowledge of the company’s tech capabilities and infrastructure from the AI conversation. The same goes for information technology (IT), with a little over a third (36%) including their IT department in AI governance. Overall, there is a lack of consensus about which roles should play a part in AI governance, with no single title — not even the new title of chief AI officer — selected by more than half of respondents, even in an ideal governance scenario.

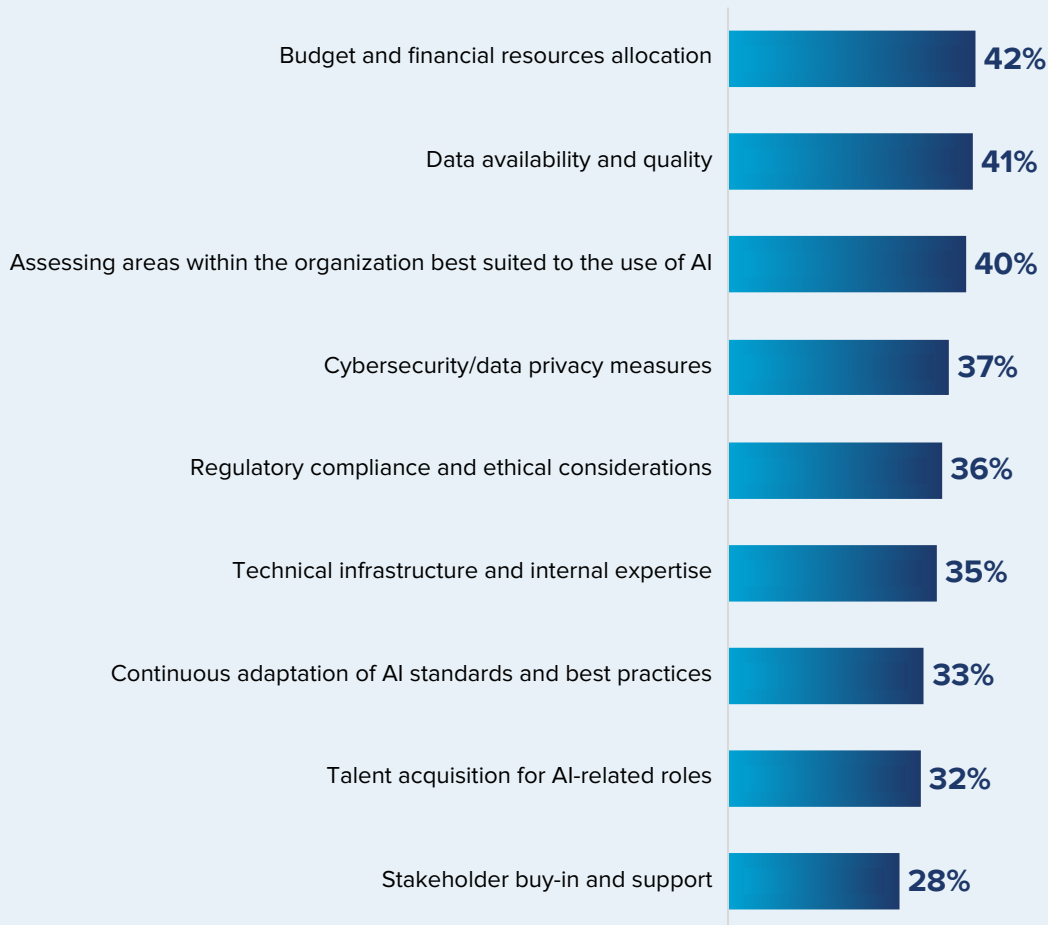
All of this is not particularly surprising — the adoption of AI is accelerating, and companies are feeling their way in terms of necessary risk management, roles, and documentation. Given governmental signaling around the risks of AI, these efforts will no doubt intensify as risk assessments are undertaken and roles and resources are allocated.



Despite the ongoing evolution in AI governance and compliance, life sciences companies are generally optimistic about their preparedness to implement AI governance and compliance, with the majority rating themselves at least somewhat prepared to address the diverse array of issues associated with implementation. Those who say they are “very” prepared, however, dips below half. Not surprisingly, given the costs and uncertainties associated with AI, four in 10 feel equipped to handle budget and financial resources allocation (42%), data availability and quality (41%), and assessing best usage areas (40%).

Lower proportions are prepared to deal with one of the most significant challenges for life sciences overall — cybersecurity (37%) — meaning that AI implementation arguably adds further stress to resources in this area. And even fewer have confidence in the internal frameworks and talent currently in place.

*Level of preparedness for addressing aspects of AI implementation
“Very prepared” responses only*



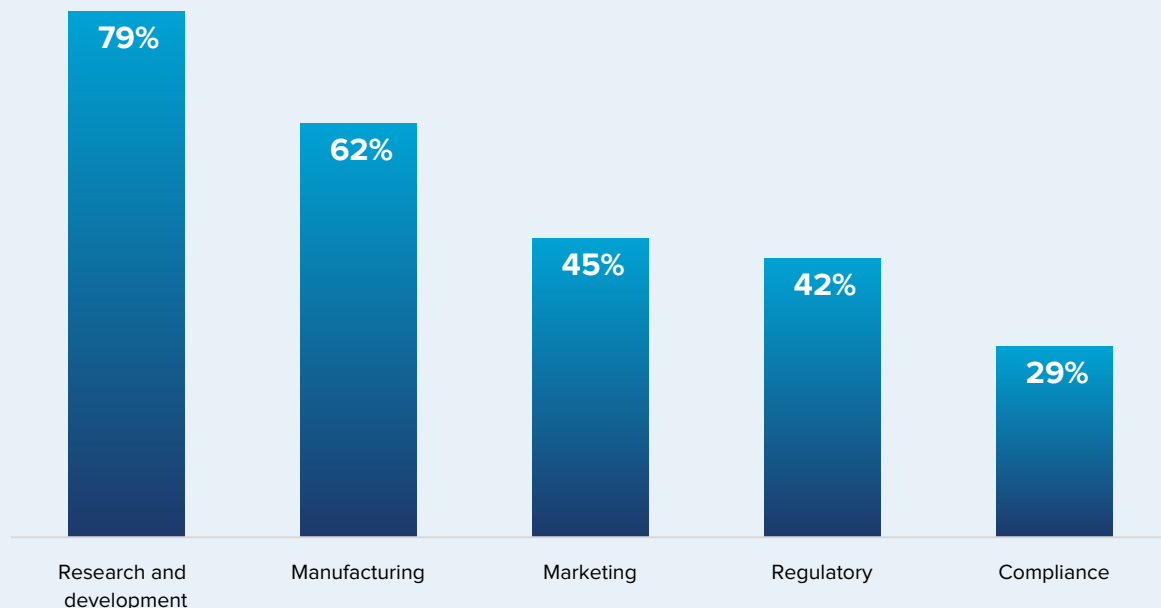
Using AI in Regulatory and Compliance Functions

Despite outstanding questions of how AI will be regulated in life sciences around the world, both companies and their regulators appear to be exploring AI applications for traditional regulatory and compliance functions, although the full impact in these areas may take time to materialize.

Of those looking to leverage AI for regulatory (42%) functions, the majority of respondents (86%) are currently using or in the process of implementing predictive AI, with generative AI tools following at 86%. Predictive AI use is even more widespread for those leveraging AI in compliance (29%), with 57% already utilizing it (compared to just 44% for regulatory). Particularly for life sciences companies working across jurisdictions with similar reporting requirements, tools that can help automate or streamline repetitive statements and identify potential gaps in compliance could both save time and offer risk mitigation advantages.




Areas where organizations have implemented or are considering implementing AI



Life sciences companies expect government agencies like the U.S. Food and Drug Administration (FDA) and the European Medicines Agency (EMA) to seize these advantages too. Nearly half (45%) of respondents believe regulators will use AI to analyze large volumes of public documents to identify misstatements in the next year, while 43% are anticipating regulators to use AI to automate routine compliance oversight tasks.

Expectation for government agencies' use of AI over the next year



Survey Conclusion

The rapid pace of AI implementation in life sciences requires action now to anticipate and manage risks. As best practices, applications, regulations, and risks evolve alongside AI technologies, life sciences companies will need to stay up to date as they continue to leverage these tools across the enterprise.

In particular, it is critical that companies ensure they have the proper teams and controls in place to safely implement new technology within the product lifecycle, and that these teams are equipped to assess the ongoing risks in tandem with the opportunities. The benefits that AI and other new technologies bring to life sciences are innumerable, but so too are the repercussions of mismatched governance, oversight, and expertise.

In the next section of this report, we will explore some of the issues raised by the use of AI and approaches to managing risk.

Understanding and Managing the Life Sciences/AI Convergence

The risks of AI are already in the sights of government enforcement agencies. As current U.S. Deputy Attorney General Monaco noted in her [remarks](#) earlier this year:

Like a firearm, AI can also enhance the danger of a crime. Going forward, where Department of Justice prosecutors can seek stiffer sentences for offenses made significantly more dangerous by the misuse of AI — they will. And if we determine that existing sentencing enhancements don't adequately address the harms caused by misuse of AI, we will seek reforms to those enhancements to close that gap. As it did with cyber, the law governing AI will develop. But our existing laws offer a firm foundation. We must remember that.

In the EU, the authorities have already set out their views on AI governance, and the EU AI Act sets out comprehensive regulation on AI. Substantial fines can be imposed for noncompliance through the AI Act, and an AI Office and AI Board have been set up to coordinate national authorities responsible for enforcing the AI Act. EU institutions are also discussing new legislation that will specifically cover liability for damage caused by AI systems, which will likely increase the potential for enforcement in this area.

Enforcement risks are present throughout the life of a product developed using AI. However,

enforcement agencies and plaintiff's lawyers will be particularly focused on the use of AI in regulatory, quality and product commercialization applications. Some potential areas of risk that have been targeted or suggested include:


- Difficulties in verifying and validating the design and controls associated with AI use, particularly in the context of regulatory inspections and scrutiny of data integrity.
- “AI-washing” via promotional or investor statements that overstate the role of AI in ensuring product safety and effectiveness, or other attributes.
- Deployment of digital tools to track patient care that could provide more than de minimis value under fraud and abuse laws, including sharing inaccurate information that could be characterized as inducing false claims or providing data that targets a healthcare practitioner or patient in ways that may be characterized as unlawful.
- AI-powered review of materials that creates promotional vulnerabilities.
- Subtle promotional placement in ambiguous “white coat” contexts without being transparent about a source or sponsorship.
- AI use in patient and reimbursement support that is excessive or excessively reduces physician responsibilities and costs.
- Commercial uses of AI that risk compromising patient privacy or are vulnerable to cyberattacks.

- AI-pricing algorithms that may be used in a manner that could be characterized as fixing prices or allocating markets.
- A failure to obtain appropriate approval or clearance for prescription drug use related software (PDURS) incorporating AI, or inaccuracies in such software that drive patient adverse events, skew clinical decision-making, or broaden patient use beyond approved labeling.

Entirely new theories and methods of enforcement will doubtless develop as AI use cases emerge that will make it even more difficult to include a “human in the loop” to ensure appropriate risk management and accountability.

The key to managing these risks will be taking a comprehensive and probing risk-based approach to ensure that the risks of various AI uses are appropriately assessed and, where necessary, mitigated. As Robert DeConti, chief counsel to the inspector general’s office in the U.S. Department of Health and Human Services, [recently stated](#):

...compliance professionals... are going to be among the very first to spot where... AI has created a program integrity vulnerability, and maybe you will be using your own AI tool to detect it. When the new technology is rolled out in your organization (and maybe it already has been), and processes are further automated, I anticipate that the compliance department will likely be on the front lines dealing with situations where the provider used AI algorithms to bill for things that did not actually happen, or services that the patient did not need.



Such risk management efforts should be built upon a strong understanding of the developing and multifaceted global frameworks for AI regulation, as well as the development of appropriate internal governance structures, policies, and procedures that are adapted over time to reflect technological developments, government policy changes, and enforcement actions.

In each AI application, companies must then focus on an analysis of the nature of the data and AI use being considered, as well as control plans for the software and the mapping of potential risks associated with the AI tool’s development and targeted use. This will require the adoption or acquisition of compliance monitoring and investigation resources — technological, financial, and human — to ensure that the industry’s important objectives can be met while also maintaining a level of control over the application of the technology that is consistent with applicable law and prevailing compliance expectations.

The following sections focus on the developing global framework for AI, issues arising from AI integration, and the mitigation of AI risks.

The Developing Framework for Global AI Regulation

Governments in Europe and North America broadly agree on the catalog of risks presented by AI but are pursuing divergent paths when it comes to addressing them — creating a complex and evolving global regulatory environment that poses challenges for life sciences companies.

To better understand what's driving the survey results, consider the developing, complex landscape that life sciences companies currently face. The U.S. has taken a largely sectoral approach to AI regulation at the federal level, applying regulators' existing statutes to new technologies, such as the Federal Trade Commission's authority to [take action](#) against companies using AI to engage in discriminatory practices or the FDA's oversight of AI tools classified as medical devices.

The Biden administration has taken small steps toward horizontal regulation, developing and applying policy frameworks across agencies' efforts. In 2023, they laid out an initial vision for the emerging AI industry in the [Executive Order on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence](#), which is built around eight primary objectives: (1) establish new standards for AI safety and security; (2) promote responsible innovation and competition; (3) support workers; (4) advance equity and civil rights; (5) protect consumers; (6) protect privacy and civil liberties; (7) use AI responsibly in the federal government and build federal AI governance capacity; and (8) advance American leadership in global AI governance. By mandating a set of minimum evaluation, monitoring,

and risk-mitigation practices for use in the federal government, the Biden administration is attempting to use the federal example and procurement policy to foster responsible AI deployment and development in the private sector as well. To that end, the executive order also called for various agencies to undertake rulemakings and other inquiries related to AI. While President-elect Donald Trump has vowed to repeal the executive order, the standards for federal government usage and procurement policies may largely survive. On the whole, however, the Trump administration seems likely to proceed even more cautiously on AI regulation than the Biden team.

Congress, and especially the Senate, undertook a crash course on AI and its implications in 2023. Drawing on months of forums, briefings, and listening sessions, the Bipartisan Senate AI Working Group released "Driving U.S. Innovation in Artificial Intelligence: A Roadmap for Artificial Intelligence Policy in the United States Senate" in May 2024, emphasizing support for innovation and offering a more limited approach to regulation.

The House of Representatives similarly established a bipartisan Task Force on Artificial Intelligence to develop a legislative framework for that chamber. In the interim, the House Republican leadership came out against prohibitions on algorithmic discrimination that had been included in the bipartisan privacy bill negotiated between the chairs of the House Energy and Commerce Committee and the Senate Commerce, Science, and Transportation Committee. Accordingly, it seems unlikely the House Task Force will propose significant regulation when it presents

its report, probably later this year. If Congress adopts a package of AI bills before the end of the year, the package likely will promote research, education, and training, and the development of voluntary best practices — but not adopt new regulation.

While the prospects for significant congressional action on regulation remain dim, U.S. state governments have been a hive of activity. Almost half of states have consumer privacy laws that regulate automated decision-making (ADM). A few state (and local) statutes regulate particular AI applications (e.g., Illinois’s law on using AI to evaluate video interviews in hiring). Colorado has enacted the first U.S. statute generally regulating AI — it targets discrimination in AI systems that make “consequential” decisions about individuals, including with respect to healthcare services. California also recently adopted laws on disclosure and detection of AI-generated content and on disclosures about the data used to train AI systems. A wide range of AI legislation remains under consideration in state capitols, with additional measures likely to be adopted over the next few years.

In contrast to the U.S.’s light-touch and sectoral approach, the EU has chosen to regulate AI intensively and horizontally, with broad measures covering the entire economy. As a starting point, AI systems using personal data must comply with the [General Data Protection Regulation \(GDPR\)](#), which has a number of provisions addressing AI development and deployment.

On top of the GDPR, the EU’s [AI Act](#), which came into force this year, takes a hybrid approach, following the template of EU product-safety legislation with a risk-based approach to regulation while also aiming to protect [fundamental rights](#).

It has a wide reach. In addition to European developers and deployers of AI systems, the AI Act applies to non-EU developers that want to market their models and systems inside the EU — and even to non-EU deployers of AI systems from which the output is sent into the EU.

A small set of practices posing an unacceptable risk to fundamental rights are prohibited outright, such as expansion of facial recognition databases from untargeted scraping of internet or CCTV footage and biometric categorization by certain sensitive or protected attributes.

The AI Act prescriptively regulates AI systems used in high-risk use cases such as safety components for various types of regulated products and — unless there is no significant risk to health, safety, or fundamental rights — various applications of biometrics; employment decisions; and access to, or eligibility for, healthcare and other essential services and benefits.

Even non-high-risk AI systems must comply with various transparency requirements. For example, systems intended to interact with people must make it obvious that they are AI, and generative AI outputs must be identified as such.

Most of these obligations fall on the provider (i.e., developer) of the AI system, not the deployer. But a deployer will be treated as the provider if it puts its name or trademark on a high-risk AI system already on market, substantially modifies an existing high-risk AI system, or modifies a non-high-risk AI system to become high risk.

Paradoxically, for a highly prescriptive piece of legislation, the AI Act is in many ways unfinished. A lot of details have been left to implementing

and delegated acts by the European Commission; guidelines and other guidance from the European Commission, EU AI Office, EU member states' authorities, and other bodies; and technical standards to be adopted by the European standards-setting bodies that will provide safe harbors for compliance by providers of high-risk AI systems and general-purpose AI models.

The EU's revised Product Liability Directive will require member states to subject AI systems and other software to their product liability laws. In addition, the EU's co-legislators are considering an AI Liability Directive to clarify how these product liability laws will apply to AI systems.

Similar to the EU, in the UK, the GDPR already governs significant aspects of the development and deployment of AI systems using personal data. But, as in the U.S., the UK is primarily taking a sectoral approach to AI regulation, with existing regulators applying high-level principles to AI use in their domains. The Starmer government seems likely to continue its predecessor's principal focus on the safety of AI frontier models.

These differences are only the tip of the iceberg, as other countries and even states within countries adopt their own AI frameworks. China has adopted a raft of measures that aggregate into relatively comprehensive AI regulation. These measures combine the consumer and worker protections common to Western AI regulation with provisions to maintain social stability and party control. A number of countries have privacy laws like the GDPR that cover ADM, and their regulators regularly provide guidance on how their privacy laws apply to the development and deployment

of AI. Brazil and Canada are among the countries seriously considering AI legislation. Japan may follow suit, as the ruling Liberal Democratic Party issued a 2023 white paper suggesting that AI regulation may be necessary. Other jurisdictions like Singapore continue to believe that only soft-law guidance on AI governance is necessary.

International regulatory consensus is unlikely to materialize any time soon, especially since it took months for just the G7 nations to agree to [high-level AI principles](#) last year.

The current landscape for privacy laws, particularly within the U.S. but elsewhere as well, illustrates how complicated the regulatory framework for AI might become for multinational companies to navigate. The rapid development of privacy statutes and regulations in the past decade has led to inconsistent requirements across jurisdictions, creating potential confusion among both businesses and consumers, and has diverted resources to compliance that companies could have invested in new products and services.

However, even if the legal regimes remain very different, international standards could help to bridge the differences. The International Organization for Standardization has already issued standards on AI Management Systems (ISO 42001) and AI Guidance on Risk Management (ISO 23894). They will not suffice for compliance with the EU AI Act, but they are a start. Major jurisdictions eventually may recognize the same standards as consistent with their own laws, which would harmonize the laws for global businesses. Until then, however, companies will have to decide how much regulatory dissonance they can handle.

Key Regulatory Considerations in Adopting AI Technologies

With all AI's promise in product development, however, one key consideration is the eventual regulatory approach to the specific use of AI, which varies by jurisdiction and the application and attributes of the technology.

United States

In the U.S., an important consideration for pharmaceutical and medical technology companies developing software that incorporates AI or ML is whether such products are subject to regulation by the FDA as medical devices. While not all AI/ML healthcare technologies fall under the FDA's purview, assessing whether such technologies are FDA-regulated requires a product-specific analysis of intended use and technological characteristics. Technologies that employ AI/ML-based software can be subject to FDA oversight as devices if the product meets the statutory definition of a medical device and does not fall within the scope of a statutory exemption or an FDA enforcement discretion policy.

Although the FDA has proposed frameworks for the regulation of AI/ML-based devices, under current law, AI/ML technologies that meet the definition of a device are regulated under the same general FDA framework for other software-based medical devices. To date, the FDA has authorized for marketing over 900 AI/ML-enabled medical devices, including many novel and innovative uses of AI/ML for patient care.

When Are AI Tools Medical Devices?

Key questions to consider when evaluating AI tools intended for health-related uses include the following:

- Does the AI tool meet the statutory definition of a device? Is the tool intended for the treatment, diagnosis, cure, or prevention of a disease or condition? Is the tool intended to affect the structure or function of the body?
- Does the AI tool fall within the scope of a 21st Century Cures Act (Cures Act) exclusion for certain low-risk software functions? The Cures Act exclusions encompass certain electronic health record software functions, administrative software functions, general wellness software functions, medical device data system software functions, and clinical decision support (CDS) software functions for healthcare professionals.

The Cures Act non-device CDS exclusion, in particular, is one that many sponsors developing novel AI/ML-enabled decision support technologies often seek to take advantage of. While these types of AI tools can potentially fall within the scope of the CDS exclusion, one of the challenges is whether the AI and its logic can be adequately explained in the CDS tool's labeling to satisfy the requirement that a non-device CDS be intended to enable a healthcare professional to independently review the basis for the software's recommendation.

- If the AI tool meets the statutory definition of a device, does the AI tool potentially fall under one or more FDA guidance enforcement discretion policies for low-risk device software functions (e.g., software functions that perform simple calculations for healthcare professionals)?
- When analyzing whether a device AI tool falls under an FDA enforcement discretion policy for low-risk software devices, FDA has cautioned that generative AI tools in particular have unique characteristics that can make it difficult to determine the bounds of a product's intended use, and therefore whether the generative AI tool is subject to enforcement discretion. For a proposed generative AI tool, part of the regulatory status analysis should include consideration of whether incorporation of the generative AI introduces potential uncertainty or risk in a product that would otherwise, without generative AI, have been considered low risk and not the focus of FDA's regulatory oversight.
- If the AI tool meets the statutory definition of a device and is not subject to enforcement discretion, what device classification applies to the tool? Does the tool require FDA marketing authorization (e.g., 510(k) clearance, premarket approval)? Is there an existing predicate device? If the AI tool is a novel, unclassified device, is it a candidate for classification through the de novo process?
- For regulated device AI tools, what other FDA regulatory controls apply to the development and commercialization of the product (e.g., investigational device exemption requirements, establishment registration and device listing, post-market safety reporting, design controls and quality system compliance, labeling)?
- Where multiple entities are involved in the development and commercialization of a regulated device AI tool, to which entity or entities do the various FDA regulatory controls apply (e.g., specification developer vs. manufacturer)?
- For regulated device AI tools, is there a process in place to analyze if and when improvements or modifications made to the model, including modifications made by the model learning from real-world data, trigger a requirement to obtain a new or supplemental FDA device marketing authorization? Should clearance or approval of a predetermined change control plan (PCCP) be considered to help facilitate such modifications to the AI tool?
- How to address and minimize potential risks with AI models and how best to apply a good machine learning practice (GMLP) during development and commercialization to help ensure the safety, quality, and reliability of the AI/ML-enabled device?

While the FDA has not issued any GMLP regulations or formal guidance, the agency has issued guiding principles in collaboration with the Medicines and Healthcare products Regulatory Agency (MHRA) in the UK and with Health Canada. These include, but are not limited to, adopting a total product lifecycle approach to the development of AI/ML models, ensuring datasets are based on a representative population and bias is minimized, implementing good software engineering and security practices, ensuring the model is tested during clinically relevant conditions, and monitoring the performance of the model over time.

Prescription Drug Use Related Software



If an AI tool's outputs relate to a drug, an additional consideration is whether the AI tool could be subject to FDA drug labeling oversight as PDURS. Notably, software, including AI/ML-enabled software, can be considered PDURS even if it does not meet the definition of a medical device. Under a proposed FDA definition, PDURS generally includes software that (1) is disseminated by or on behalf of a drug sponsor and (2) produces an end-user output that supplements, explains, or is otherwise textually related to one or more of the sponsor's drug products.

Notably, the end-user outputs of PDURS can potentially be considered promotional labeling or FDA-required labeling for the underlying drug and be subject to FDA drug labeling requirements. For PDURS outputs considered promotional labeling, this includes a requirement to submit the outputs to the FDA through the Form 2253 process. Under FDA draft guidance, the specifics of the drug labeling-related FDA requirements that apply to PDURS outputs depend in part on whether the PDURS is also subject to review by the FDA as a device.

Europe

Similar to the U.S., in the EU, AI/ML systems that meet the definition of a medical device may fall under the regulatory framework for devices, as set out in the EU Medical Devices Regulation (MDR) for general medical devices and the EU In Vitro Diagnostic Medical Devices Regulation (IVDR) for in vitro diagnostic products. Determining whether an AI/ML system falls within the scope of the MDR and/or IVDR requires an analysis of whether the intended purpose of the technology has a "medical purpose" as set out in the legislation, including being intended to diagnose, prevent, monitor, predict, treat, or alleviate a disease. This determination is based on both the functionality of the technology and the claims made about its intended uses. Further, under the MDR and IVDR, software medical devices, including AI/ML technologies, are subject to conformity assessment with the assistance of a third-party notified body before being placed on the market, meaning the process can be time-consuming.

The significant divergence in the EU compared to many other jurisdictions is that the EU has adopted so-called “horizontal” legislation: the EU AI Act. This has broad applications to AI systems across different sectors and uses, including those that fall within the definition of a medical device. Starting August 2, 2027, the AI Act will impose a new regulatory system on high-risk AI systems, which will include most AI/ML-based medical devices, in addition to the regulations that already exist with the MDR and IVDR. There are provisions in the AI Act that attempt to reduce the amount of duplication between the regimes and it is expected that specific guidance on how the AI Act applies to medical devices will be prepared. However, it is not yet clear to what extent any such standards will acknowledge the specific context of AI/ML-based devices and whether they will also help demonstrate conformity with medical device legislation.

The EMA has also recently finalized its “Reflection Paper” on how it will approach the output from AI/ML systems within the assessment process when an AI/ML system is used in the lifecycle of biopharmaceutical products, and published guidance on the use of large language models in regulatory science and the regulation of medicinal products. The EMA has made clear that the use of AI should comply with existing rules on data requirements as applicable to the particular function that the AI is undertaking. It is clear that any data generated by AI/ML will be closely scrutinized by the EMA, and a risk-based approach should be taken depending on the AI functionality and the use for which the data is generated.

In the UK, the government previously indicated that it does not intend to introduce “horizontal” legislation regulating AI with broad applications across sectors — in contrast to the AI Act. Rather, the UK’s approach has been to set broad guiding principles to be implemented by the relevant sector regulator as part of the existing regulatory frameworks. Accordingly, AI/ML-based devices are regulated under the legislation applicable to medical devices. In addition, the MHRA has developed guidance, standards, and policies specific to AI/ML-based devices. These include guiding principles on PCCPs and GMLP (as noted above, in collaboration with the U.S. and Canada), as well as establishing a regulatory sandbox, the “AI Airlock,” for the testing of AI/ML-based devices. The UK’s new Labour government has so far continued with many of the initiatives that had already been introduced, and not indicated any major deviations from the previous approach other than some limited measures that are unlikely to be relevant to AI/ML-based devices.

Questions to Ask Before Deploying AI



As with any program, planning and preparation are key to successful AI implementation. Life sciences companies looking to use AI tools should first ask themselves:

- 1**
Do we have an AI policy and governance structure in place?
- 2**
Do we have a system for conducting thorough risk assessments when procuring AI tools, including diligence on both the vendor and the tool?
- 3**
Have we implemented robust data privacy and security measures?
- 4**
Have we cleaned any datasets to be used in training the AI system?
- 5**
Have we updated our procurement templates to reflect AI-specific challenges?
- 6**
Have we trained both procurement personnel and end users on the use of AI tools?
- 7**
Is there any need or requirement to tell customers, stakeholders, or investors about our proposed use of AI tools?

AI Privacy and Cybersecurity Risks

Large datasets are the essential building blocks of AI — presenting inherent privacy and cybersecurity concerns for life sciences companies looking to leverage new technology.

The use of personal data (including personal health data) for training AI models requires prior notice and often consent. However, this can be difficult to obtain when datasets come from clinical trials or treatment where the future use of AI is typically a secondary purpose. And our old (preexisting) datasets were obtained when no one thought to give notice or ask for consent. Failures of notice and consent can lead to data or algorithmic disgorgement. But deidentifying or anonymizing personal health data, although complex, can alleviate some privacy concerns in the use of the health datasets.



AI can also introduce new avenues of attack into organizations by hackers, through vulnerabilities in the technology and in the training datasets. This can be particularly problematic when dealing with sensitive health data, given the potentially steep fines following data breaches in both the U.S. and EU.

To help mitigate these risks across the product lifecycle, life sciences companies should consider:

- Entering into contractual provisions with companies providing datasets for AI training or the AI tool itself to clarify the privacy and cybersecurity requirements vendors have in place.
- Future-proofing data for AI training purposes by ensuring that notice and consents for clinical trials and in treatment settings contain provisions encompassing that use.
- Creating internal policies and procedures about the appropriate use of personal data in AI.
- Vetting AI use cases that include privacy and cybersecurity considerations with cross-functional teams with relevant experience.
- Considering cybersecurity assessments and testing for vulnerabilities both *before* and *after* implementing AI.

IP Risks in AI and Life Sciences



The growing use of AI throughout the product lifecycle also presents significant IP hurdles for the life sciences field, particularly around questions involving patentability for AI-assisted drug development. It's an issue the industry is bracing for — nearly all (97%) of respondents report some degree of concern about facing AI-related IP issues in the next year, with the majority (74%) reporting a high level of concern.

One likely cause for concern: courts and most patent offices around the world have consistently required an inventor to be a natural person. The U.S. even goes a step further, mandating humans make a [significant contribution](#) to the invention.

Inventorship of patents claiming AI-assisted therapeutic inventions therefore must be considered carefully, as a lack of proper

inventorship could be used to challenge patents during litigation. When engaging attorneys to help with drafting patents for products whose development involved the use of AI tools, life sciences companies should ensure that they're aware of the associated risks and comply with applicable rules in each of the relevant patent offices.

When it comes to life sciences products, this requires more than simply taking the output of the AI tool and claiming that output as their own invention. Examples of significant contributions could include structural modification of the AI tool's drug output, identifying the methodology to create the modified drug, or conducting the synthesis of the modified drug. Designing, building, or training an AI system to solve a specific problem could also be sufficient for joint inventorship in certain cases.

However, under this framework, a researcher who merely provides inputs or prompts to an existing AI system to identify drug candidates is typically *not* considered a co-inventor. This is because they are not contributing with a specific problem in mind or to obtain a particular output to solve a problem — they’re simply exercising normal skills expected for the field.

Additionally, when considering patent eligibility, U.S. applicants can claim AI-powered solutions that involve improvements to the subject matter, such as a method of treatment, but not those that merely improve the AI model itself.

Similarly, the European Patent Office has issued rulings on inventions asserting to be autonomously generated by AI, holding they did not meet the criteria of an inventor as it was not a natural person. In the UK, the ability of AI systems to be named as the inventors in patent applications was addressed in 2023 by the UK Supreme Court in *Thaler v. Comptroller-General of Patents, Designs and Trade Marks*. This was consistent with the position in the U.S.; the UK Supreme Court held that a natural person must be the inventor, and accordingly, there is no right in the UK conferred on an applicant

to obtain a patent for a product or process created or generated by an AI system. The UK Intellectual Property Office updated its Manual of Patent Practice in April 2024 following the Supreme Court decision to confirm that an inventor must be a natural person.

Users must also consider the AI-driven patent-drafting risks. The U.S. Patent and Trademark Office published guidance on the use of AI, in which it acknowledged that AI may be used to prepare and prosecute patent applications but cautioned on the risks of the use of AI, such as incomplete or inaccurate outputs and the sharing of sensitive and confidential client information to third-party AI systems, which may implicate national security, export control, and foreign filing license issues. The guidance emphasizes that the duty of candor and good faith extends not only to the personal actions taken by individuals associated with a case, but also to the actions these individuals take with any AI tools. The users must thoroughly review the AI’s output for completeness and accuracy of AI draft applications, claims, and information disclosure statements prior to filing. The European Patent Office and the UK Intellectual Property Office have yet to publish equivalent guidance.

Promotion of AI Use and Avoiding “AI Washing”

As life sciences companies incorporate AI into their business, some may seek to ensure that customer and investors are aware of these efforts. While accurate statements about AI use are certainly acceptable in many contexts, the U.S. Securities and Exchange Commission (SEC) has indicated it will take aggressive enforcement action against companies making false claims about AI, also known as “AI washing.” Similar to its greenwashing campaign, the SEC appears to be scrutinizing the accuracy and robustness of companies’ disclosures about the use of AI — with a laser focus on actual or potential false or misleading statements — and is expected to further develop AI-related disclosure requirements.

Private securities litigation, which frequently targets life sciences companies, is also beginning to focus on false and misleading statements made about AI. It is only a matter of time until life sciences companies become the targets of these claims.

The challenge for life sciences companies at this juncture is to properly identify and disclose how they are using AI, while simultaneously building a governance structure and disclosure controls around its use to mitigate SEC enforcement and civil litigation risks, even as the technology is rapidly changing. This requires the expertise and dexterity to understand the emerging technology and to develop a flexible framework that can adapt as technology advances.

To mitigate enforcement and civil litigation risks, life sciences companies should:

- Revisit the existing SEC regulatory framework and how it may apply to AI. Advancements in AI will likely outpace regulatory developments, which may bring increased challenges as older regulations are applied to new technologies.
- Understand how AI technology is being used within the organization, including how the company relies on it, how it may produce unintended results, how it may not suit the purpose it is trying to achieve, and how different AI products may interact with each other. This is especially critical for those in management, legal, and compliance roles.
- Build the right governance structure and internal controls around the use of AI upfront, including designing a principled approach to AI’s use and reliance on that use, as well as management oversight over the disclosure process, understanding that AI technology may advance faster than the protocols.
- Consider company-specific disclosures rather than boilerplate language when disclosing material risks related to AI. Define what the company means by “AI,” including how and where it is being used by the company and whether the relevant technology is developed in-house or supplied by others.
- Make sure that policies, procedures, and disclosure controls address AI usage and provide checks and balances to ensure validity and consistency of messaging across AI disclosures.

AI in Litigation



Litigation attorneys at life sciences companies have long heard about generative AI's potential to revolutionize the handling of cases. However, navigating the hype around generative AI and assessing its risks can be challenging, and associated risks must be managed.

Use Cases for Generative AI in Litigation

Generative AI tools can replace or enhance traditional legal tasks, including legal writing and research, e-discovery, and early case assessment. For example, generative AI may be helpful in creating rote, basic, or template-style documents, such as correspondence, pleadings, briefs, or deposition outlines, or for testing alternative phrasing as drafts are refined. It can also search databases of legal authority quickly and summarize relevant authority with citations. However, there have been high-profile cases of attorneys filing AI-generated documents with errors. And while legal technology companies

are constantly improving their generative AI tools, issues persist. AI might simply make something up (or “hallucinate”), fail to correct a mistaken assumption in a prompt, or rely on outdated authority. Attorneys must critically verify generative AI outputs by reviewing and understanding sources and independently assessing their authority and relevance.

E-discovery

Attorneys have also long used AI to manage the staggering amounts of data in modern litigation practice. Generative AI can assist with reviewing and synthesizing large volumes of discovery, targeting appropriate data for collection and analysis, and generating privilege logs. Early pilot programs for document review indicate that the tools can meet or exceed the accuracy of human review and previous technology solutions, but they can be costly and require significant “machine time” and iterative assessment and vetting.



Early Case Assessment

When documents relevant to a case are input into a generative AI tool, the tool can provide attorneys with information to help them assess the strength of claims or defenses early on. Through active engagement with the generative AI model — inputting prompts and asking questions — attorneys can identify key actors or significant facts faster than with traditional early case assessment tools like concept clustering or relationship analyses.

Managing Risks

Generative AI tools also introduce various risks in litigation that must be considered, including costs, data protection, and ethical concerns. First, the costs associated with generative AI tools are often high, with subscription and usage fees alone being prohibitive. These tools break down submitted content into smaller fragments (known as “tokens”) and charge fees for each token fragment, increasing costs dramatically for

larger volumes of data. Additionally, ensuring the tools’ outputs are accurate and sufficient can be time-consuming, and attorneys and technology provider engineers may need to develop custom solutions to meet the required goals, which can take up considerable resources.

Privacy, security, and ethical issues abound in this area. Generative AI tools might expose client information to unauthorized access if not properly secured, and sharing data with third-party AI providers can risk violating confidentiality agreements and data protection laws. While company attorneys don’t need to be AI experts, they must understand generative AI tools’ purposes and ensure they function properly. Vetting these tools is essential; attorneys must demand high standards, diligently assess their accuracy and effectiveness, discuss proper validation with technology providers, and ensure thorough training for case teams.

Getting Paid for AI-Related Services and Tools

When AI is a product or an aspect of a product, such as PDURS or digital diagnostic tools, it is vital to successfully navigate coverage and reimbursement. Further complicating the issue, requirements for payment policies differ across jurisdictions.

United States

In the U.S., the Centers for Medicare and Medicaid Services (CMS) has been considering an AI payment policy, including AI-related services and those that involve software algorithms, under the Medicare Physician Fee Schedule (MPFS) for a number of years, but its thinking is still evolving.

The inherent challenges to payment for AI-related services under the MPFS are driven by (1) the unique aspects of these innovative services (e.g., subscription costs, software analysis, service may augment or replace physician work) and (2) the underlying physician fee schedule (PFS) payment methodology, which is relatively resource-based and takes into account physician work (time and intensity) and practice expenses (PE) (the resources involved in furnishing services).

One key issue is that PE includes both direct (clinical labor, medical supplies, and medical equipment attributable to a patient encounter) and indirect (administrative labor, office expenses, and all other expenses) costs. Software is difficult to attribute to a particular patient or direct PE and has thus been considered an indirect cost tied to costs for associated hardware that is considered medical equipment. Moreover, the PE methodology relies on the Physician Practice Information Survey, which is conducted periodically (last done in 2007 and 2008) and does not adequately reflect emerging technologies like AI, and the next survey is not expected until 2027.

The handful of AI-related services paid for under the MPFS are instructive and reflect the CMS's attempt to recognize the use and incurred costs by practitioners for such technologies. For example, remote retinal imaging (RRI) is a diagnostic test for diabetic retinopathy that uses a software algorithm. The CMS initially contractor-priced it, then used a crosswalk methodology (moved it to another service/code under the PFS) to reflect the overall relative resource use of RRI.

The CMS has an evolving payment policy for hospital services that considers some AI-based services as software as a service (SaaS). The CMS proposed a definition for SaaS in the 2023 outpatient prospective payment system as “algorithm-driven services that assist practitioners in making clinical assessments, and that providers pay for either on a subscription or per-use basis.” However, the CMS also said that “Due to the novel and evolving nature of these technologies, it has been challenging to

compare some SaaS procedures to existing medical services for purposes of determining clinical and resource similarity.”

Payment for Digital Therapeutics

Very few digital therapeutics (DTX) have been covered under Medicare, and it remains an evolving area. A significant challenge for DTX is an available benefit category. The Medicare statute neither expressly covers digital health technologies like DTX, nor explicitly excludes them; however, such categories do not easily lend themselves to these new digital technologies.

New Technology Add-On Payments for AI-Related Technologies and Digital Diagnostics

Medicare provides temporary add-on payments to hospitals under the inpatient prospective payment system (IPPS) for certain new technologies that meet specified criteria. The CMS has also adopted alternative pathways for new technology add-on payments (NTAP), including for certain technologies that are part of the FDA’s Breakthrough Devices Program. In recent years, the CMS has awarded NTAP for a growing number of products, including a computer-aided diagnostic software to assess and characterize brain tissue abnormalities using computed tomography-guided image data.

Europe

The payment for and reimbursement of healthcare and associated treatments in the EU is governed mainly by the national laws of the EU member states; it is substantially influenced by the fact that most healthcare is socialized. This reflects one of the fundamental premises that EU member states are responsible for their national health policies: the organization, management, and delivery of health services and medical care, and the related allocation of resources, including payments and reimbursement.

There is some EU-level harmonization between the national systems, including in relation to public procurement, cross-border healthcare, and, to a limited extent, the transparency of pricing and reimbursement of medicinal products.

In general, there are well-established national procedures for determining the pricing of medical products and their inclusion in national health systems, generally based on some form of health technology assessment. However, approaches to the pricing and reimbursement of medical devices are disparate and often opaque.

In January 2025, the EU Health Technology Assessment Regulation will start to create some harmonization in relation to health technology assessment. However, national procedures will continue to be determinative of market access even when EU clinical assessments are relied upon.

Broadly speaking, there are four general commercial models for digital health services and tools in the EU member states:

1**Direct-to-consumer**

(subscription paid for by patient or provided free to patients)

2**Value-based contracting**

(provided as part of a predetermined, agreed-upon supply of a medical product)

3**Device-like reimbursement**

(paid for by authorities/insurance companies when prescribed by doctor)

4**Drug-like reimbursement**

(paid for by authorities/insurance companies when prescribed by doctor)

For example, the German healthcare system, which uses a health insurance model, has the most developed pathways to reimbursement of AI technologies in the EU. The German Digital Healthcare Act (DiGA) and associated rules came into effect in December 2019; it permits doctors to prescribe digital healthcare apps listed in the DiGA directory and have them be reimbursed by statutory health insurance. Overall, the German fast-track process has set the standard for other EU countries.

In the UK, national rules still largely reflect those applicable in the EU, particularly in relation to public procurement and transparency. All AI-based medical devices may be reimbursed in principle once the Conformité Européenne/UK Conformity Assessed (CE/UKCA) mark has been granted. However, in practice, National Health Service (NHS) access will be available only when it is recommended by the National Institute for Health and Care Excellence (NICE) and the product meets the Digital Technology Assessment Criteria.

NICE has developed a specific program, the Medical Technologies Evaluation Programme, to support the assessment of digital health technologies. It evaluates medical devices to determine whether they are either cost-saving or cost-neutral, and provides guidance to the NHS. However, such guidance is not mandatory and does not automatically result in funding by local NHS bodies.

While the UK government has established an innovative pathway that provides access to innovative medical devices, including AI, in circumstances where there is no single national reimbursement pathway, access to these technologies is inconsistent.

Mitigating Risk in Negotiating AI Vendor Contracts



Life sciences companies should keep in mind the following critical issues when negotiating in-license and services agreements with AI vendors:

1

Understand the Terms

AI tools may be offered under different terms — “free,” “consumer,” or “enterprise” — with varying levels of protection for the customer. Life sciences companies should review any online policies (such as terms of service, application programming interface terms, acceptable use policies, customer support policies, privacy policies, and data processing agreements) and make sure the terms they negotiate with the vendor control the agreement.

Additionally, some AI tools are built on third-party foundation models not controlled by the

vendor that may be subject to separate terms and conditions. Life sciences companies should explore how this may impact their use of the tool.

2

Identify the Source of Training Data

To date, many developers have used datasets scraped from the internet to train generative AI applications. This information could contain errors or biases, so it’s important to understand the inherent flaws or limitations in such tools’ training datasets, particularly if the tool will be used for important functions like R&D. What’s more, scraped datasets have been the subject of dozens of IP lawsuits, which could lead to injunctions prohibiting the operation of, or even requiring the destruction of, those AI tools.

3

Require Monitoring and Compliance

While it's always critical to require IT vendors to monitor and comply with applicable statutes, regulations, orders, and other laws, it's particularly important for AI vendors given the speed at which legislatures and regulators are adopting AI-specific statutes and regulations at all levels of government.

4

Mandate Cooperation in Audits and Investigations

Life sciences companies should require their AI vendors to inform them of any audits or investigations, provide any information and assistance necessary to respond, and coordinate any of the vendor's responses, to the extent permitted by law.

5

Ensure Rights to AI Outputs

Life sciences companies should be clear as to what rights they — and the vendor — have to AI outputs. At a minimum, customers should ensure that they have sufficient rights to use the output for any intended uses, whether or not they “own” the outputs.

6

Restrict Vendor Control

AI vendors, particularly those selling generative AI tools, often attempt to obtain broad rights to use the prompts that their customers enter into the AI tools, including for improvements, training, or providing services to other customers. However, life sciences companies' AI prompts may contain company trade secrets, protected health information, or other sensitive data. Therefore, it may be appropriate — or even legally required — to prohibit or restrict the vendor from storing and using those prompts for any purpose other than to provide the necessary service to the customer.

Even when a vendor has rights to use customer prompts for other purposes, it will still be important to maintain confidentiality and restrict vendors from allowing tools to incorporate portions of those prompts in future outputs for other customers.

7

Look for Warranties, Indemnities, and Remedies

AI vendors should provide robust representations and warranties regarding the accuracy and reliability of their tools and the vendors' right and ability to provide their tools. Those should be backed up by adequate indemnities and other remedies in the event of a breach.

Glossary

Artificial intelligence (AI)

AI leverages technology to emulate human intelligence in performing tasks. It has a broad range of applications in life sciences, ranging from product discovery, diagnostics, and virtual assistants to predictive analytics tools for personalized medicine and enhanced product commercialization.

Generative AI

Generative AI refers to algorithms or neural networks capable of identifying patterns and structures within data and generating diverse content types, from audio to text to images.

Machine learning (ML)

ML is a subset of AI and the foundation for both predictive and generative AI. It enables the development of intelligent systems by allowing algorithms to learn from vast and complex datasets without explicit programming, thereby facilitating tasks and advancements.

Predictive AI

Predictive AI employs statistical algorithms to forecast future events or trends by analyzing historical data patterns, thus facilitating informed decision-making and proactive interventions to optimize processes.

Security by design

Security by design involves integrating security measures into the initial design and development stages of a product, system, or application, prioritizing the inherent inclusion of security features rather than retroactively adding them.


Methodology

Arnold & Porter surveyed 100 senior executives and department heads specializing in AI, risk management, legal counsel, data science, quality assurance, regulatory affairs, compliance, security, marketing, and digital strategy with the help of third-party B2B panel providers. The online survey ran from May 2024-June 2024.

The respondents, who played a leading or supporting role in decision-making around AI implementation, represented various sectors within the life sciences and healthcare industries. Responses are anonymous and data were analyzed in the aggregate.

About Arnold & Porter

Arnold & Porter combines sophisticated regulatory, litigation, and transactional capabilities to resolve clients' most complex issues. With over 1,000 lawyers practicing in 15 offices worldwide and 200 lawyers in the life sciences field, we offer deep industry experience and an integrated approach that spans more than 40 practice areas. Through multidisciplinary collaboration and focused industry experience, we provide innovative and effective solutions to mitigate risks, address challenges, and achieve successful outcomes.



Acknowledgements

Arnold & Porter commissioned the research and survey analysis for this report from [Greentarget Global Group](#). The report benefited from the contributions, insights, and review of a diverse team of Arnold & Porter lawyers, including:

Daniel A. Kracov

Chair, Global Life Sciences Industry Group

Veronica E. Callahan

Chair, Securities Enforcement & Litigation

Mahnu V. Davar

Co-Chair, Life Sciences & Healthcare Regulatory

Abeba Habtemariam

Partner, Life Sciences & Healthcare Regulatory

Daniel M. Hawke

Partner, Securities Enforcement & Litigation

Thomas Magnani

Partner, Corporate & Finance

Jacqueline Mulryne

Partner, Life Sciences & Healthcare Regulatory

Jane Norberg

Partner, Securities Enforcement & Litigation

Alexander Roussanov

Partner, Life Sciences & Healthcare Regulatory

Howard Sklamberg

Partner, Life Sciences & Healthcare Regulatory

Abigail Struthers

Partner, Intellectual Property & Technology

Jami Vibbert

Chair, Privacy, Cybersecurity & Data Strategy

Peter J. Schildkraut

Senior Counsel, Telecommunications

John McInnes

M.D., J.D., Counsel, Life Sciences & Healthcare Regulatory

Monique Nolan

M.D., J.D., Counsel, Life Sciences & Healthcare Regulatory

Melissa L. Weberman

Counsel, eData

Alice Ho

Ph.D, Senior Associate, Intellectual Property & Technology

Arnold & Porter's project team would like to thank those that helped in the review and production process, providing clarity and context to the report, including Justin Antonipillai, founder and CEO of WireWheel.

In addition, the project team extends sincere thanks to the experts who provided their time and insights, as well as the 100 senior industry executives and department heads who participated in the online survey.

The views and opinions expressed in this report represent those of Arnold & Porter and the aggregate results of the survey and expert interviews, and do not necessarily reflect those of any specific individual or organization mentioned above.

Arnold & Porter

Questions about this report or Arnold & Porter?

Please email us at AISurvey@arnoldporter.com

