



European Federation of Pharmaceutical  
Industries and Associations



# AI Across the Medicines Lifecycle

Insights from Preliminary Case Studies  
and Considerations for Policy

November 2025

# Executive summary

**Artificial Intelligence (AI) is transforming the pharmaceutical landscape, offering unprecedented opportunities across the medicines lifecycle, from drug discovery and clinical development to manufacturing and post-approval safety monitoring. As AI becomes more deeply embedded in regulatory decision-making and processes that impact patients, the need for robust governance becomes critical. This report presents a cross-industry perspective on how AI governance can be integrated effectively into medicinal product research and development (R&D) and post-authorisation settings, with a focus on ensuring trust, transparency, and regulatory alignment through ongoing dialogue between industry and regulators.**

## Addressing regulatory and ethical needs

Regulatory agencies, including the European Medicines Agency (EMA), have emphasised the need for trustworthy and explainable AI (e.g., EMA 2023 Reflection paper on the use of Artificial Intelligence (AI) in medicinal product lifecycle<sup>1</sup>). However, while current regulatory guidance establishes a foundation for governing AI, additional clarity is needed in how AI should be governed throughout its lifecycle. This report offers real-world insights from pharmaceutical leaders, based on interviews structured around AI risk and governance grounded in international best practices across industries (e.g., Organisation for Economic Co-operation and Development (OECD), National Institute of Standards and Technology (NIST), International Organization for Standardization (ISO) and International Electrotechnical Commission (IEC), European Union (EU) Agency for Cybersecurity (ENISA), and others). The report aims to articulate opportunities and challenges which could be used by regulators to draft further guidance on AI.

## Methodology and purpose

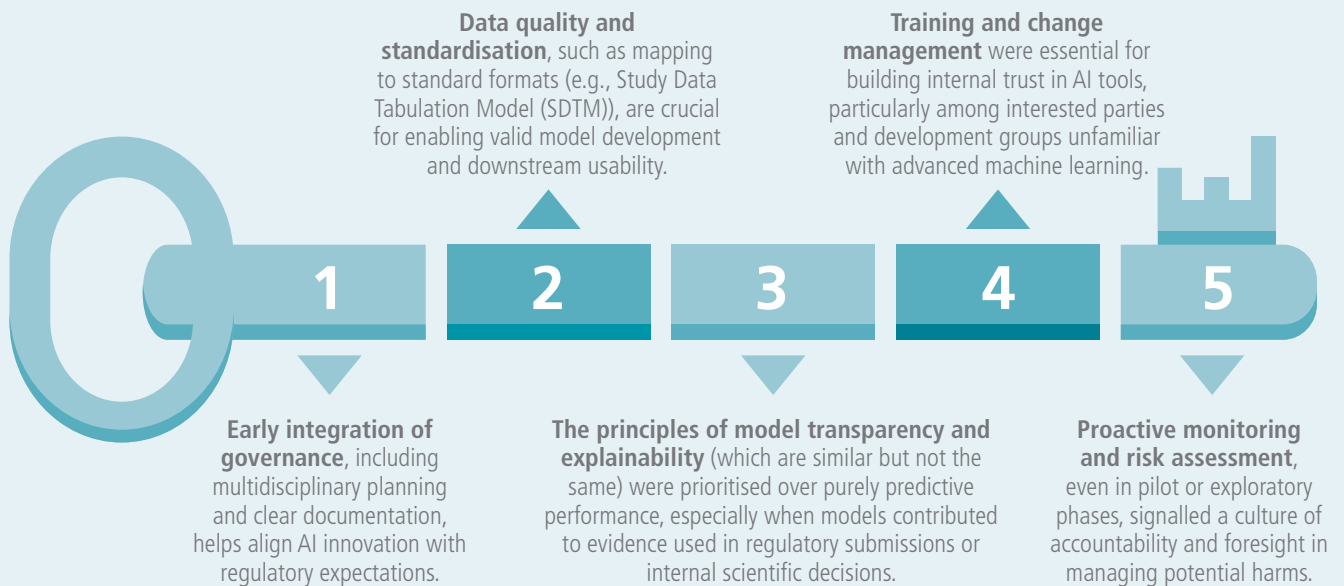
The report draws on a series of in-depth interviews with executives and technical experts across major pharmaceutical companies. Each of the four case studies explores a distinct AI application in the medicines lifecycle, examining both the technical approach and the governance structures in place. To ensure comparability and relevance, a standardised interview guide based on an internationally recognised AI lifecycle was used to explore governance at each stage: planning and design, data collection, model development and validation, deployment, and ongoing monitoring.

<sup>1</sup> [https://www.ema.europa.eu/en/documents/scientific-guideline/draft-reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle\\_en.pdf](https://www.ema.europa.eu/en/documents/scientific-guideline/draft-reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf)



## Key insights and good practices

Across the case studies, several common themes emerged around best practices:



These practices demonstrate that high-performing AI systems can be developed responsibly, with appropriate checks and balances across the lifecycle. These insights capture what companies report doing today in their AI governance practice. The policy recommendations that follow reflect what interviewees believe regulators could do next, i.e., they are complementary to, rather than directly derived from, these findings.

### Recommendations to policymakers

Derived from interview questions on shaping AI policy and best practices, the recommendations below complement the governance findings. To support continued responsible innovation, the report recommends that the EMA, European Commission, and other regulators:

- Clarify the scope and application of AI-related exemptions (e.g., the scientific R&D exemption in the EU AI Act, including the possibility to conduct clinical trials by putting products into service prior to CE marking).
- Encourage iterative dialogue and opportunities for engagement between regulators and industry on policy themes for AI, with transparency through greater sharing of AI use cases and governance strategies.
- Promote dynamic and adaptive guidance formats, which are more easy and quicker to update, such as Q&A/FAQ type documents.

- Support regulatory efforts to collaborate and harmonise where possible on AI expectations, and consider the use of regulatory sandboxes to explore industry-led AI use cases and governance solutions.
- Integrate AI oversight into existing regulatory frameworks, such as endpoint qualification/validation for AI-based imaging and RWE processes for AI-generated synthetic controls, to ensure consistency, efficiency, and trust.

AI holds the potential to revolutionise medicines development, enabling smarter, more efficient development processes. However, realising this potential requires trust, both in the technology and in the processes that govern it. Embedding governance across the AI lifecycle can help pharmaceutical companies ensure the use of AI is effective while also being ethically aligned with evolving regulatory standards. This report serves as a foundation for further dialogue between industry and policymakers on appropriate regulatory and governance frameworks, based on industry experience, to ensure AI is developed and used in the public interest.

# Table of Contents

<b>Introduction</b>	<b>6</b>
<b>Why integrate AI governance in the medicines lifecycle?</b>	<b>6</b>
<b>Best practice AI governance model</b>	<b>7</b>
<b>Governance across the AI lifecycle: Risks and best practices</b>	<b>8</b>
1. Planning and design: Foundations for responsible innovation	8
2. Data collection and processing: Addressing bias and privacy	8
3. Model development and validation: Ensuring robustness and explainability	8
4. Deployment and use: Managing integration and adoption	8
5. Ongoing monitoring and risk mitigation: Sustaining trust over time	8
<b>Preliminary case studies: Capturing AI governance practices</b>	<b>9</b>
<b>Governing AI responsibly in clinical research: Lessons from a real-world use case</b>	<b>12</b>
The challenge: Limited insight from traditional pathology	12
The AI Solution: Unlocking hidden signals from histology slides	12
Governance in practice: Building trust through embedded oversight	12
Recommendations for policymakers	13
Summary	13
<b>Unlocking Clinical Innovation: The Promise of Synthetic Data</b>	<b>14</b>
The challenge: Data gaps in clinical trials and patient access	14
The AI solution: Synthetic data for clinical research acceleration	14
Governance in practice: Embedding trust by design	14
Recommendations for policymakers	15
Summary	15
<b>AI Governance in Pharmaceutical R&amp;D: A Path Forward for Patient Safety and Regulatory Excellence</b>	<b>16</b>
The challenge: Limitations of “traditional” quality assurance and auditing in pharmaceutical R&D	16
The AI solution: Changing R&D Quality Oversight	16
Governance in practice: AI from Concept to Patient Impact	16
Recommendations for policymakers	17
Summary	17
<b>AI in pharmacovigilance: Governance lessons from post-market surveillance</b>	<b>18</b>
The challenge: An overwhelming literature burden in drug safety	18
The AI solution: Co-developing an AE screening tool	18
Governance in practice: Applying AI principles across the lifecycle	18
Recommendations for policymakers	19
Summary	19

<b>Conclusion</b>	<b>20</b>
<b>ANNEX</b>	<b>21</b>
<b>Illustrating effective AI governance</b>	<b>21</b>
Non-clinical research	21
Early clinical research	21
Late-stage clinical trials	21
Manufacturing	21
Regulatory submission	21
Post-approval activities	21
<b>Interview guide: Capturing best practices in AI governance across the AI lifecycle</b>	<b>22</b>
Plan & design: Setting the foundation for trustworthy AI	22
Data collection & processing: Ensuring high-quality, bias-free inputs	22
Model development & testing: Ensuring performance & reliability	23
Deployment & use: Scaling AI while managing compliance	23
Ongoing monitoring & risk mitigation: Continuous oversight of AI	24
Closing: Shaping AI policy & best practices	24



## Introduction

**The integration of Artificial Intelligence (AI) across the medicines lifecycle presents both transformative opportunities and the need for clear guidance and governance. As AI systems move from computer science research laboratories into regulated settings – where they support trial design, manufacturing quality, pharmacovigilance, and even regulatory submissions – their influence over patient safety, public trust, and regulatory decision-making becomes more pronounced.**

Unlike AI in standalone digital products, AI is increasingly embedded within core pharmaceutical processes, which are already governed by mature frameworks such as good clinical practice (GCP), good manufacturing practice (GMP), good pharmacovigilance practice (GVP), and a host of ICH guidances. For regulators like the European Medicines Agency (EMA), the core challenge is how to ensure that AI use is safe, effective, and accountable.

This report aims to examine how AI is being used across the medicines lifecycle and what governance approaches industry is already applying. Our early findings, drawn from selected case studies provided by leading pharmaceutical companies, suggest that many uses of AI are already sufficiently governed under existing regulatory frameworks (e.g., GxP), and that AI-specific controls and validations are being implemented.

## Why integrate AI governance in the medicines lifecycle?

**AI systems can significantly enhance pharmaceutical research, development, and post-approval activities. However, their rapid integration into processes across the medicines lifecycle may introduce unique risks and uncertainties. Ensuring these technologies operate safely, effectively, and ethically requires a structured governance framework aligned with regulatory expectations and industry best practices.**

To frame this evolving conversation, we examined use cases along two interlocking axes:

- The medicines lifecycle: From early research to post-market surveillance, AI use tends to grow more impactful to patients, and higher risk, the closer it gets to decisions affecting patient health and regulatory filings.
- The AI risk lifecycle: Effective governance requires attention to where AI is used and how it is developed, validated, deployed, and monitored. These stages echo and align with principles from industry and public sector frameworks.

These two lifecycles are interdependent. A seemingly low-risk model that operates without basic lifecycle controls (e.g., undocumented datasets or poor validation) can still present significant risks. Conversely, a seemingly high-risk model with appropriate controls can operate safely within existing frameworks.

# Best practice AI governance model

In this context, what we mean by an AI governance model is a stage based management system for AI in medicines R&D and post authorisation activities that specifies who does what, when, and with what evidence across five stages: planning & design; data collection & processing; model development & validation; deployment & use; and ongoing monitoring & risk mitigation. It draws on recognised frameworks (e.g., OECD, NIST, ISO/IEC, IEC, ENISA) and is intended to be proportionate to inherent risk and context of use, ensuring transparency, accountability, and regulatory readiness throughout the lifecycle.

Effective AI governance in medicines is not about creating entirely new regulatory pathways, but about embedding AI oversight into established frameworks that already govern clinical research, medical devices, and real-world evidence (RWE). For example, AI-based imaging tools used to determine disease burden should be integrated into existing endpoint qualification and validation processes, ensuring that AI outputs meet the same standards of reliability and clinical relevance as traditional endpoints. Similarly, the use of AI-generated synthetic control arms in clinical trials should be governed within current RWE processes, with appropriate documentation, validation, and transparency to support regulatory decision-making. By aligning AI governance with these established processes, regulators and industry can leverage existing expertise, streamline compliance, and foster innovation without compromising safety or scientific rigor. A summary of the AI governance model is provided in Table 1.

**Table 1: AI governance model at a glance**

Lifecycle state	Purpose	Minimum controls (examples)
<b>Planning &amp; design</b>	Align intent with ethics, regulation, and patient safety	Define intended use and context of use; perform inherent risk tiering; assign roles & accountability; consult Legal, Regulatory, DPO as applicable
<b>Data collection &amp; processing</b>	Ensure lawful, high-quality, representative data	Lineage and provenance documented; quality checks; harmonisation to standards (e.g., SDTM where relevant); lawful basis, minimisation, de-identification, anonymisation; bias checks
<b>Model development &amp; validation</b>	Build robust, reproducible, transparent models	Pre-specified protocol; train, validate, test splits; independent validation; acceptance criteria; transparency and explainability appropriate to use; version control
<b>Deployment &amp; use</b>	Integrate safely into workflows with oversight	Change control and release process; user training; human-in-the-loop where required; access and security controls; incident and issue reporting
<b>Ongoing monitoring</b>	Detect drift, manage residual risk, keep evidence current	Monitoring plan with thresholds; periodic re-validation; audit trail; trigger and escalation criteria; retirement and sunset criteria

DPO = data protection officer; SDTM = study data tabulation model

This governance model provides clarity and consistency, enabling systematic identification, management, and mitigation of risks at each stage of the medicines lifecycle. It helps policymakers and industry stakeholders understand where potential issues emerge, how to intervene effectively, and how to maintain accountability throughout AI application. The narrative that follows expands each stage with practices and examples. The controls and listed in Table 1 are only some of the expected behaviours to support transparency and proportional oversight.

# Governance across the AI lifecycle: Risks and best practices

To support ethical, trustworthy and effective use of AI in the pharmaceutical sector, governance must span the entire AI lifecycle. Based on international frameworks for risk management across the AI lifecycle, this section synthesises best practices and common risks across five key phases: planning, data collection, model development and validation, deployment, and ongoing monitoring. This framing was used in structured interviews with industry leaders across pharmaceutical companies to capture the case studies in this report.

## 1. Planning and design: Foundations for responsible innovation

At the earliest stage, aligning AI development with ethical principles, regulatory requirements, and patient safety goals is essential. Companies establish cross-functional governance structures involving legal, regulatory, medical, bioethics, data science, quality management, and compliance teams. Early vetting of AI use cases against EU AI Act criteria, GDPR requirements, best practice (e.g., GxP), and emerging EMA guidance helps define clear boundaries and avoid downstream risks. Key decisions are documented, and high-risk systems are flagged for enhanced oversight.

**Areas of focus:** Inadequate alignment with regulatory expectations; failure to anticipate downstream harms; siloed decision-making.

## 2. Data collection and processing: Addressing bias and privacy

AI systems in pharma rely on complex data ecosystems (including real-world data, clinical trials, and operational data needed to demonstrate regulatory compliance) which pose challenges for harmonisation and quality control. Organisations mitigate bias using fairness audits, representative data sampling, and privacy-enhancing techniques such as synthetic data or federated learning. Ensuring alignment with GDPR includes, for example, data minimisation, anonymisation, and explicit documentation of data provenance and quality.

**Areas of focus:** Hidden biases in training data; non-compliance with GDPR; lack of data traceability or documentation.

## 3. Model development and validation: Ensuring robustness and explainability

Model development balances performance with explainability, often using techniques like SHAP, LIME, or interpretable algorithms (e.g., decision trees). Independent teams validate model performance and generalisability, applying traditional statistical methods and predefined utility benchmarks. Documenting model assumptions and testing methods supports reproducibility and facilitates regulatory review.

**Areas of focus:** Overfitting or lack of generalisability; insufficient documentation to support transparency and regulatory review; opaque models in high-risk use cases.

## 4. Deployment and use: Managing integration and adoption

Before deployment, AI tools are evaluated for regulatory readiness and integration into clinical or regulatory workflows. Companies train end users (including medical affairs, researchers, and regulatory teams) to ensure proper understanding and usage. Stakeholder communication is critical to build trust, especially for AI systems supporting or informing regulatory submissions or patient-related decisions.

**Areas of focus:** Resistance from users; incomplete integration with workflows; lack of transparency in system outputs.

## 5. Ongoing monitoring and risk mitigation: Sustaining trust over time

Post-deployment, AI systems require continuous monitoring for performance drift, new risks, or regulatory updates. Teams conduct scheduled revalidations, track emerging failures, and update documentation for audit readiness. Internal governance mechanisms evolve to respond to new EMA or EU AI Act requirements, ensuring ongoing alignment and responsiveness.

**Areas of focus:** Drift in, or collapse of, model performance; lack of version control or audit trails; insufficient agility in responding to regulatory changes.

## Preliminary case studies: Capturing AI governance practices


To move beyond theory and showcase the current state of AI governance in pharmaceutical research, a series of structured interviews were conducted with senior leaders from pharmaceutical companies. These conversations aimed to capture concrete examples of how AI is being applied across the medicines lifecycle, from research and clinical development to manufacturing and commercialisation, and how governance practices are evolving to manage associated risks. The interviews focused on the capabilities and benefits of AI while exploring the guardrails, validation protocols, and organisational structures that companies are putting in place to ensure AI is used safely, effectively, and in alignment with ethical and regulatory expectations.

The case studies that follow offer an inside view of how leading pharmaceutical companies are translating governance principles into practice. By aligning the interviews with a structured AI lifecycle model, adapted from international best practice frameworks such as OECD, NIST, ISO/IEC, ENISA and others, these examples illustrate common patterns in AI governance, highlight persistent challenges specific to implementing AI responsibly in pharmaceutical settings, and showcase emerging solutions developed to address these governance challenges. Each case highlights different use cases and stages in the AI lifecycle, collectively offering a grounded perspective on what best practice looks like today and what policy gaps remain.


In each case study, we distinguish between challenges encountered in traditional approaches and new challenges that arise specifically from the use and governance of AI. The solutions described focus on how organisations have addressed these AI governance challenges in practice. Policy recommendations are provided to help regulators support these solutions and overcome persistent barriers identified in the use cases. A summary is provided in Table 2.



Table 2: Preliminary case studies, challenges, governance responses, and linked policy asks

 <b>Case study: Governing AI responsibly in clinical research: Lessons from a real-world use case</b>	
<b>Traditional challenge</b>	Pathology workflows compress rich slide data into single scores (e.g., TPS), limiting insight and consistency.
<b>AI governance challenge (what governance had to solve)</b>	Clarify applicability of R&D exemption for retrospective, non-interventional analysis; assure transparency vs. explainability balance; ensure data provenance and representativeness and address drift between commercial and trial samples; manage third-party and vendor transparency; define proportionate oversight and monitoring in a retrospective setting.
<b>Solution in the use case</b>	Multidisciplinary governance panel and accountability reports; pathologist-in-the-loop reviews; model cards and data sheets; fine-tuning to handle data drift; qualitative monitoring; scope kept retrospective (no patient-care impact).
<b>Policy recommendations linked to the use case</b>	Clarify R&D exemption; adopt risk-based, proportional oversight; prioritise transparency (documentation and auditability) over abstract explainability; clarify EMA/EC roles; co-develop best practices with industry.

 <b>Case study: Unlocking Clinical Innovation: The Promise of Synthetic Data</b>	
<b>Traditional challenge</b>	Trial enrolment and RWD gaps slow studies; reliance on limited or siloed datasets.
<b>AI governance challenge (what governance had to solve)</b>	Third-party vetting and GDPR alignment; robust lineage, quality and harmonisation (e.g., SDTM); fairness and bias assessment; balance interpretability vs. performance; build AI literacy and trust among clinical stakeholders; define monitoring for responsible use.
<b>Solution in the use case</b>	Multidisciplinary governance (Regulatory/DPO/data science); partner vetting + internal checklists; lineage and quality controls and fairness metrics; utility and performance benchmarks; preference for interpretable approaches where effective (e.g., sequential trees); training and workshops; iterative reviews.
<b>Policy recommendations linked to the use case</b>	Encourage transparency & knowledge-sharing; adopt stepwise guidance (terminology to principles to risk framework to living Q&A); clarify exemptions; strengthen industry–regulator engagement; promote international harmonisation (EMA–FDA–ICH).



## Case study: AI Governance in Pharmaceutical R&D: A Path Forward for Patient Safety and Regulatory Excellence (R&D Quality oversight)

### Traditional challenge

Traditional QA and auditing is resource-intensive and retrospective; late detection of issues (e.g., AE under-reporting).

### AI governance challenge (what governance had to solve)

Data harmonisation and appropriate anonymisation; ensure reliability and reproducibility (GMLP); choose explainable options when performance is comparable; align to computer system validation SOPs with clear intended use; drive adoption & data literacy; post-deployment monitoring and change control.

### Solution in the use case

AI ethics principles and checklist; collaboration with data management for lineage and anonymisation; apply GMLP; prefer explainable models when feasible; maintain an AI inventory; follow CSV tied to intended use; open-source release via IMPALA; change-management and training; periodic updates and reviews.

### Policy recommendations linked to the use case

Clarify definitions (e.g., transparency vs. explainability); build on existing GCP and GVP frameworks; context-sensitive, risk-tiered regulation; enable deeper regulatory dialogue with industry.



## Case study: AI in pharmacovigilance: Governance lessons from post-market surveillance (literature screening)

### Traditional challenge

Manual screening of vast literature with very low AE yield; heavy, repetitive workload to maintain compliance and traceability.

### AI governance challenge (what governance had to solve)

Licensing and IP constraints for text & data mining; limited explainability vs. need for outcome validation; define deployment boundaries with QA; unclear AI-literacy expectations; ensure organisational readiness and oversight.

### Solution in the use case

Early legal and compliance review and cross-functional risk assessment; structured documentation and quality assurance; independent validation plus peer-review and vendor testing; pilot with comparison to manual review; training to build AI literacy; gated rollout with oversight.

### Policy recommendations linked to the use case

Issue domain-specific guidance (e.g., EMA on GenAI for PV); create early engagement pathways (cf. FDA EDSTP); publish precedents; clarify GxP interpretation for AI; clarify TDM and licensing rights for safety use.

**AE** Adverse Event  
**AI** Artificial Intelligence  
**CSV** Comma Separated Value  
**DPO** Data Protection Office  
**EC** European Commission  
**EDSTP** Emerging Drug Safety Technology Program  
**EMA** European Medicines Agency  
**FDA** Food and Drug Administration  
**GDPR** General Data Protection Regulation

**GenAI** Generative AI  
**GCP** Good Clinical Practice  
**GMLP** Good Machine Learning Practice  
**GVP** Good Pharmacovigilance Practice  
**GxP** Good Practice  
**ICH** International Council of Harmonisation  
**IMPALA** Inter Company Quality Analytics  
**IP** Intellectual Property  
**PV** Pharmacovigilance

**QA** Quality Assurance  
**Q&A** Question & Answer  
**RWD** Real-World Data  
**R&D** Research & Development  
**SDTM** Study Data Tabulation Model  
**SOP** Standard Operating Procedure  
**TPS** Tumour Proportion Score

# Governing AI responsibly in clinical research: Lessons from a real-world use case



## Summary

**AI is transforming biomedical research through credible, well-governed innovations like this one. By extracting richer features from pathology images, this system supports biomarker discovery and enhances scientific understanding in a way that is robust, interpretable, and low risk.**

But innovation is only one dimension. Trust is earned through governance that is proportionate, transparent, and well-aligned with regulatory expectations. This case shows that it's possible, and that the life sciences industry is ready to help shape a path forward.

With the right regulatory scaffolding, AI can deliver insight and impact. Policymakers have the opportunity and responsibility to ensure that future AI applications in research are ambitious and accountable.

## The challenge: Limited insight from traditional pathology

Clinical research depends on a deep understanding of how patients respond to therapies. Biomarker exploration is a critical part of this process—it reveals biological mechanisms, informs trial design, and guides therapeutic decisions. Yet traditional pathology methods offer limited resolution: in standard workflows, pathologists assign a single score (such as Tumour Proportion Score, or TPS) to each slide, compressing rich histological data into a single metric.

The opportunity to extract more insight from these slides is clear—but doing so reliably, ethically, and transparently is where the challenge lies.

## The AI Solution: Unlocking hidden signals from histology slides

In a recent initiative, an AI-powered pipeline was developed to analyse IHC-stained tumour histology slides from a Phase 2 clinical trial. The goal was to generate deeper, spatially informed features that could serve as exploratory biomarkers. Unlike conventional pathology scoring, the AI system identifies specific tissue structures, quantifies cell types, and produces outputs that are both richer and more consistent.

Importantly, this AI system was used for retrospective, non-interventional analysis. It did not influence treatment decisions. It supported scientific exploration by helping researchers better understand the tumour microenvironment and its relationship to patient outcomes.

The value is clear: more nuanced biological signals, uncovered from existing data, with the potential to inform future trial strategies and therapeutic development.



## Governance in practice: Building trust through embedded oversight

While the AI system offered clear scientific potential, its development was shaped by a strong governance framework that balanced innovation with accountability. The system was reviewed under a structured internal governance model, combining risk-based assessment, transparent documentation, and ethical oversight. Below, we trace how governance was applied across key phases of the AI lifecycle.

### Planning and risk assessment

Although the project began before the formal governance process was fully established, it was retrospectively registered and reviewed. Oversight was provided by a Responsible AI team and a multidisciplinary governance panel, composed of experts in ethics, regulation, and technical development.

Each AI initiative within the organisation must submit an accountability report that outlines data provenance, intended use, model documentation, risk categorisation, and oversight plans. For this use case, the system was classified as low risk—based on its retrospective nature, non-clinical use, and distance from patient decision-making. The concept of proximity to patient, i.e. how many defensive layers are there between the patients and the AI system, should be a central criterion in tiering AI risk.

### Data acquisition and preparation

The AI models were trained on commercially sourced pathology images that were carefully selected to mirror the clinical trial population by matching factors such as cancer type, disease stage, and the medical procedure used to obtain the tissue sample. This alignment ensured that the AI system would be relevant to the scientific questions under study.

Procurement followed established risk protocols, with third-party vendors subject to the same standards as internal teams. Any limitations in transparency from external partners were addressed through internal validation of data integrity and model robustness.

### Model development and validation

Development was deeply collaborative. Pathologists reviewed model predictions, helping interpret overlays and identify potential issues. This partnership served two purposes: enabling clinical validation of model behaviour and enhancing transparency around what the AI system was actually “seeing.”

Good machine learning practices were followed throughout, including training-validation-test splits and the use of holdout datasets. Where data drift was identified between commercial and clinical samples, the models were fine-tuned accordingly.

The AI system was not continuously deployed or retrained. Its purpose was to support a fixed analysis within a defined study period, allowing governance practices to focus on traceability and documentation rather than real-time monitoring.

### Explainability and transparency

The model’s outputs, such as cell-type distributions and whole-slide TPS, were directly interpretable by clinical experts. This reduced the need for abstract explainability tools. Instead, emphasis was placed on transparency by including detailed documentation through model cards, data sheets, and impact summaries to make the system auditable and understandable.

The company advocates for this shift in regulatory thinking—from explainability as a technical ideal to transparency as a practical standard, especially in use cases with well-understood outputs.

### Monitoring and accountability

Even in retrospective settings, monitoring practices were clearly defined. Model performance was qualitatively reviewed by pathologists, with annotations used to detect drift or inconsistencies. Poorly modelled samples were excluded, and all results were recorded in updated accountability reports, which were subject to audit by both engineering and policy reviewers.

## Recommendations for policymakers

This use case highlights how structured AI governance can be applied in practice, and what regulators might consider as they refine policy frameworks. Several lessons emerge:

- Clarify R&D exemption: Retrospective, exploratory research like this project appears to fall under the R&D exemption in the EU AI Act. Clear, binding guidance could reduce uncertainty and ensure continued innovation in non-interventional settings.
- Support proportional, risk-based oversight: Regulatory requirements could scale with the AI system’s proximity to clinical decision-making. A low-risk, retrospective research tool should not be subject to the same burden as an AI used in direct patient care. Tiered regulation could improve both compliance and innovation.
- Prioritise transparency over explainability: For many domain-specific applications, especially in biomedical research, outputs are interpretable without the need for abstract post hoc explanations. Policymakers could prioritise transparency (e.g., documentation, auditability) as the preferred foundation of trust while ensuring protection of proprietary algorithms, training data and trade secrets to maintain innovation incentives.
- Continue with agency collaboration and defining of roles: The relationship between the EMA and the European Commission on AI governance is still evolving. Coordinated guidance, clear responsibilities, and regulatory coherence could provide greater certainty for developers and reviewers alike.
- Collaborate on industry best practices: Many life sciences organisations are already implementing rigorous, science-aligned governance models. Policymakers could actively collaborate with these early adopters to shape practical, evidence-based regulation.

# Unlocking Clinical Innovation: The Promise of Synthetic Data



## Summary

**The adoption of AI-driven synthetic data generation holds promise to overcome potential challenges in certain clinical research settings. Synthetic data is an enabler, that can integrate, support and amplify traditional tools. It provides an alternative to anonymisation for data sharing, thus improving privacy and data utility, which can be utilised for collaboration and knowledge generation. Synthetic data solutions can also have intrinsic value beyond data sharing such as the use conditional generation to augment sparse datasets.**

Robust AI governance frameworks, emphasising transparency, ethical integrity, and continuous oversight, are essential to harnessing these benefits responsibly.

Policymakers play a critical role by providing clear, flexible, and harmonised regulatory frameworks, enabling the pharmaceutical industry to innovate confidently and efficiently. Collaborative, proactive regulatory engagement will undoubtedly foster a landscape where AI technologies can flourish, significantly benefiting patient care and enhancing clinical research outcomes.

## The challenge: Data gaps in clinical trials and patient access

Clinical trials may encounter challenges such as limited patient enrolment and retention, especially for heterogeneous chronic or rare diseases, and insufficient availability of high-quality real-world data (RWD) which may expose patients to unproven therapies until results are evident. These constraints delay research and for patients, this can mean slower access to new, potentially life-saving therapies and continued uncertainty in treatment options. Addressing these challenges can rely on approaches that can complement traditional methods without compromising scientific integrity or patient privacy and safety.

## The AI solution: Synthetic data for clinical research acceleration

The generation of synthetic data through the use of AI offers a powerful solution by creating highly realistic, fully artificial datasets created in digital worlds rather than collected from or measured in the real world. This technology enables researchers to fill certain data gaps while preserving privacy and protecting sensitive patient information. In this use case, AI-generated synthetic data was used to simulate external control arms for clinical trials, with the purpose of understanding the technology and its limitations and providing a pathway to more inclusive and efficient studies. This approach could accelerate research timelines and drive innovation in drug development and clinical research. Besides, it allows the exploration of “what if” scenarios, increases trial feasibility and strengthen evidence and assumption made based on the data from single arms of the trial.



## Governance in practice: Embedding trust by design

The use of AI, especially in sensitive domains like clinical trials, necessitates rigorous governance to uphold patient safety, ethical standards, and regulatory compliance. Comprehensive AI governance frameworks ensure careful oversight across all stages of AI application, addressing risks such as bias, privacy breaches, and data inaccuracies. By embedding transparency and accountability into AI processes, governance frameworks enhance trust among stakeholders and enable responsible innovation, paving the way for sustained advancements in medical research.

### Building foundations

Effective governance begins at the initial planning stage. This project implemented a multidisciplinary governance structure, bringing together experts from data science, medical teams, regulatory compliance, and the Data Protection Officer (DPO). Rigorous vetting of third-party partners for alignment with GDPR and EU AI Act was a core step, complemented by an internal checklist ensuring systematic adherence to evolving regulatory landscapes.

### Ensuring quality and fairness

Harmonising clinical trial data and electronic medical records (EMRs) into standardised formats such as the Study Data Tabulation Model (SDTM) is crucial for data interoperability and regulatory readiness. The team meticulously tracked data lineage and conducted iterative validation to maintain high data quality. To proactively tackle potential biases, standardised statistical fairness metrics were applied, protecting against demographic disparities and ensuring robust, equitable dataset representation.

### Balancing performance and interpretability

Multiple AI approaches were explored, including sequential trees, Bayesian networks, and transformer-based models to identify optimal performance while balancing interpretability. Sequential trees emerged as particularly effective, offering clarity and robust outcomes. Utility metrics and traditional statistical benchmarks validated the accuracy and reliability of the synthetic data, reinforcing confidence in its applicability for clinical research.

### Bridging scepticism and innovation

While this particular use case remains exploratory, significant efforts were dedicated to internal training and education, emphasising the technology's capabilities and boundaries, with the goal of fostering understanding and trust among clinical researchers and statisticians, some of whom could be traditionally wary of AI methods. Initial workshops and clear, open communication channels helped clarify the practical limits and potentials of synthetic data use, bridging gaps between AI experts and clinical teams.

### Continuous vigilance

Even though real-time deployment was out of scope of this initial exploration, proactive measures ensured continuous evaluation of model reliability and performance. Iterative qualitative assessments and internal stakeholder reviews effectively mitigated risks related to model drift, emerging regulatory standards, and trust erosion.

## Recommendations for policymakers

As the pharmaceutical industry pioneers AI applications like its use to augment synthetic data generation, regulators are playing an important role in guiding responsible innovation. Insights from this use case point to areas where thoughtful policy updates could reduce uncertainty, align global standards, and accelerate responsible adoption of AI while safeguarding patient interest.

- Enhance transparency and knowledge-sharing: Encourage greater sharing of AI use cases, including successes, challenges, and practical insights, supporting faster adoption of best practices and more informed policymaking.
- Adopt a stepwise approach to regulation: Starting at terminology definition and alignment, to overarching general principles, to risk assessment framework to living Q&A documents, rather than static guidelines, to effectively address the rapidly evolving AI technology landscape.
- Clarify regulatory exemptions: Precisely define the scope and limitations of the R&D exemption within the EU AI Act to reduce uncertainty and encourage responsible AI innovation in pharmaceutical research.
- Continue to strengthen industry collaboration: Facilitate regular dialogue and establish dedicated industry-regulator workshops to foster continuous mutual understanding and timely feedback on emerging AI governance issues.
- Promote international harmonisation: Prioritise collaboration and alignment between EMA, FDA, and other global regulators through international forums like the International Council for Harmonisation (ICH) to avoid fragmented regulatory environments.

# AI Governance in Pharmaceutical R&D: A Path Forward for Patient Safety and Regulatory Excellence



## Summary

**AI governance in pharmaceutical R&D is an enabler of innovation and patient safety. By embedding robust top-down governance practices into every phase of AI development and use, companies can significantly enhance patient outcomes and regulatory confidence. Policymakers have a critical role to play in creating clear, flexible, and context-sensitive regulatory frameworks. Collaborative efforts between industry and regulatory bodies will be essential to ensuring AI's transformative potential is fully realised, benefiting patients and the broader healthcare ecosystem alike.**


## The challenge: Limitations of “traditional” quality assurance and auditing in pharmaceutical R&D

Compliance with Good Clinical Practice (GCP) and Good Pharmacovigilance Practices (GVP) ensures the rights, safety, and well-being of research subjects and patients while maintaining the reliability of clinical research and pharmacovigilance data. Traditional quality assurance practices, such as audits and on-site visits, are resource-intensive, retrospective and can lead to late detected/undetected issues.

AI and Advanced Analytics enable R&D Quality organisations to transition from these resource intensive methods to generating more transparent, near-real-time, and holistic insights. These insights support robust decision-making and quality conclusions, enhancing the overall efficiency and effectiveness of audits and risk assessments, and ensuring compliance with ICH-GCP and EU GVP. For example, accurate and timely reporting of adverse events (AEs) in clinical trials is crucial to ensuring data integrity and patient safety. However, AE under-reporting remains a challenge, often highlighted in Good Clinical Practice (GCP) audits and inspections. Traditional detection methods, such as on-site investigator audits via manual source data verification (SDV), have limitations.

## The AI solution: Changing R&D Quality Oversight

Addressing this, the open-source R package {simaerep} was developed to facilitate rapid, comprehensive, and near-real-time detection of AE under-reporting at each clinical trial site. This package leverages patient-level AE and visit data for its analyses. The open-source package can be embedded into audits to enable fast, holistic, and repeatable quality oversight of clinical trials. The statistical probability of a site under-reporting adverse events can be used to manage, target and focus quality assurance activities. It follows general good practices and standards for R packages and has a high unit test coverage which is tested by an automated pipeline which creates a validation report that is attached to the latest release.



## Governance in practice: AI from Concept to Patient Impact

### Embedding Ethical Principles from the Start

The global Bioethics team in the company developed AI ethics principles which should be considered in the early planning of any AI-driven tools, including transparency, explainability, and human-in-the-loop. To facilitate this process, tools (e.g. AI Ethics Checklist) and a global Ethics consult can be leveraged by the development teams.

When designing a new AI system or tool, a cost-benefit analysis is completed where cost estimations of developing, piloting, and rolling-out to production are mapped against the benefits. This analysis is revisited at every project milestone and is a helpful prioritisation tool.

### Safeguarding Data Integrity and Privacy

In R&D quality oversight, harmonisation of different data sources, e.g., clinical data from eCRFs and operational data from the clinical trial management system, is challenging and typically requires a step of data cleaning and restructuring before the analysis. To ensure a clear data lineage and the right level of data anonymisation, the team works with the data management team. To ensure reliability, reproducibility and generalisability of results during model development and validation, the AI ethics principles, together with Good Machine Learning Practices are followed. In the event of “difficult to explain” models (aka “black boxes” AI models) it is always a trade-off between performance and explainability, but for a similar level of performance the most explainable and sustainable model will be chosen as guided by the AI ethics principles.

### Ensuring Transparency in Model Development

After the tool has gone through proof of concept and pilot stages, it is entered into an inventory of every AI use case in the company

and is ready to enter the regulated space. The roll-out process is guided by a computer system validation procedure, which includes all the documentation and checklists that the developers must fill in and ensure they comply with. For that process, it is important that the intended use (i.e. context of use) of the AI tools is clearly defined, as they will drive the requirements as described in the company’s AI/ML Computer System Validation Standard Operating Procedures (SOPs).

As an additional step to transparency, the {simaerep} package was released as an open-source package through the Inter coMPany quALity Analytics (IMPALA) consortium. Releasing as open-source builds confidence among regulators and stakeholders. The IMPALA consortium, established in October 2022, brings together biopharmaceutical sponsors to share knowledge and explore opportunities in advanced analytics and AI for R&D Quality.

### Facilitating User Adoption and Data Literacy

Adoption of AI solutions requires targeted change management and data literacy from the users, as well as an understanding of the intended use of the solutions. An effective approach to increase adoption has been to capture AI use for the particular use case in the users’ goals.

### Continuous Monitoring and Model Improvement

Post-deployment governance includes rigorous documentation and periodic updates to AI models, such as replacement of models when newer, more efficient ones have been developed. Regular reviews ensure continuous improvement in the available models and compliance, maintaining reliability, and safeguarding patient interests.

## Recommendations for policymakers

To further strengthen AI governance in pharmaceutical R&D, policymakers should focus on several key actions:

- Clarify essential definitions: establish clear, universal definitions of related but distinct terms like explainability and transparency across regulatory bodies.
- Build on existing regulation: regulatory requirements should build on existing frameworks, such as GCP or GVP requirements.
- Contextualise AI regulations: tailor regulations to differentiate between high-risk autonomous AI applications and lower-risk, human-supervised tools typically employed in pharmaceutical R&D.
- Encourage meaningful regulatory dialogue: foster regular, substantive exchanges between regulators and industry experts, enabling deeper mutual understanding of practical AI applications and use cases and their implications.

# AI in pharmacovigilance: Governance lessons from post-market surveillance



## Summary

**AI offers a credible path to scaling pharmacovigilance efforts—if implemented with governance rigor. This case shows how cross-functional collaboration, attention to legal and ethical risks, and clear validation protocols can enable deployment of AI that is ethical, trustworthy and impactful.**

To unlock these benefits more widely, EMA could prioritise issuing practical, use-case-aligned guidance and invest in mechanisms for early engagement. Responsible AI in safety surveillance is achievable and urgently needed to protect patients and to allow for generation of reliable evidence in an era of diverse data.

## The challenge: An overwhelming literature burden in drug safety

Post-market surveillance is essential to ensure patient safety once medicines are authorised and in use. Regulatory frameworks require pharmaceutical companies to monitor the scientific literature for adverse events (AEs) and product-related safety signals. Yet as the volume of medical publications continues to grow, the proportion containing relevant AE information remains extremely low. This imbalance creates a heavy manual workload for pharmacovigilance teams—who must review vast quantities of text to find the few that matter.

Safety reviewers spend hundreds of hours annually on literature screening. The goal was clear: maintain compliance with safety reporting requirements while seeking a more efficient, scalable solution.

## The AI solution: Co-developing an AE screening tool

To address this, a tool, previously co-developed with an external vendor to assist as safety net in the identification of potential adverse events in text, was used. The tool uses trained classifiers to decide if there is a drug of interest being mentioned in text and if there is an adverse event reported.

In this literature screening use case, as opposed to the previous “safety net” use cases, this deterministic AI makes the decision of whether there’s a marketed drug and an adverse event, rather than confirming a decision already made by a human. Early pilot testing showed the potential to significantly reduce manual burden while improving consistency and traceability.

Notably, the machine learning model that powers the tool is a GXP validated model originally used for handling small medical information requests in a safety net capacity (i.e. after a human decision was made). The team leveraged this previous work as a regulatory foundation, enabling a faster and more structured deployment of the new pharmacovigilance medical literature screening tool.



## Governance in practice: AI from Concept to Patient Impact

The governance process for this tool was intentionally robust and aligned with recognised AI lifecycle stages—from planning to deployment.

### Planning and legal review

Governance started early on. Legal and compliance teams reviewed intellectual property restrictions, particularly licensing constraints on medical literature. The team flagged early that medical journal content may not legally allow text mining—even for safety purposes—under standard license terms. Ethical, cybersecurity, and regulatory reviews followed, with a cross-functional risk assessment defining development conditions and deployment boundaries.

### Development and validation

The initial ML model was trained using labelled data from medical affairs and pharmacovigilance professionals, with clear documentation of training performance and output consistency. Internal reviewers validated its performance against expert human judgment. While explainability was limited due to model design, outcome validation was a more meaningful form of assurance

than model interpretability alone.

The performance of the AI model was documented not only internally but also through peer-reviewed publications and vendor-led external testing—adding credibility to the model’s safety, accuracy, and reproducibility.

### Deployment with oversight

Following structured documentation and quality assurance checks, the tool advanced to pilot use within the company’s local literature screening process. The end users, who currently do the task manually, compared the AI tool’s output to their manual screening prospectively and reported time savings without loss of quality.

### Organisational readiness and skills

Recognising the importance of AI literacy, the company invested in training aligned with expected requirements under the EU AI Act. While over half of staff now use internal AI tools regularly (both, generative AI and deterministic AI), the team noted a lack of clear guidance on what constitutes sufficient AI literacy in the EU AI Act for compliance purposes.

## Recommendations for policymakers

This case study demonstrates that AI can be responsibly implemented in pharmacovigilance with measurable efficiency gains while maintaining patient safety. However, policy gaps remain:

- Produce domain-specific guidance: EMA publishing guidance on the use of GenAI for PV purposes, covering key aspects such as sample sizes, validation procedures, extent of human oversight etc. would support innovations in the field.
- Introduce an early engagement pathway: Initiatives such as the FDA’s Emerging Drug Safety Technology Program offers an iterative, dialog-driven engagement model that helps build trust during AI development<sup>3</sup>. Similar early engagement mechanisms at EMA would be highly beneficial.
- Provide transparency on precedents: Publication of use cases by the regulators in partnership with Trade

associations would push the whole field forward. When regulators openly share detailed use cases and precedents, companies gain concrete examples of what kinds of AI applications are acceptable, and under what conditions.

- Clarify GxP interpretation for AI: There remains uncertainty about how GxP principles (e.g., validation, auditability) should be interpreted in practice for AI-based tools. This case shows how performance validation and audit trails can serve as effective substitutes for model explainability, but clearer policy statements would provide confidence.
- Consider licensing for safety use: European policy should consider clarifying rights around text and data mining for safety reporting. If medical content is subject to restrictive licenses, even well-governed AI solutions may not be deployable.

<sup>3</sup> The principle of model transparency is making model operations and outputs clear and auditable; the principle of model explainability is making model decisions understandable to users.

## Conclusion

**AI is moving beyond experimental pilot studies in pharmaceutical development; AI is becoming operational. Integrating structured AI governance within the medicines lifecycle offers a practical and transparent framework to support regulatory oversight and foster responsible innovation.**

These preliminary case studies show that meaningful governance is already happening across the medicines lifecycle, particularly when risk and impact are clearly understood. They illustrate successful governance practices, providing a valuable foundation for ongoing dialogue between industry and regulators, ultimately enhancing patient safety and trust in pharmaceutical innovations.

As AI applications become more embedded in medicines development, it is essential that regulatory processes for AI are integrated into the frameworks that already govern endpoints, devices, and real-world evidence. This approach enables regulators and sponsors to apply familiar standards and procedures to novel technologies, facilitating adoption while maintaining robust oversight.

As EMA and other stakeholders explore the regulatory future of AI in medicines, we recommend a continued focus on:

- Promoting transparency, auditability and knowledge sharing across all tiers of AI use.
- Supporting harmonisation and integration of AI governance into existing regulatory frameworks (e.g., endpoint qualification/validation, device processes, RWE processes).
- Recognising domain expertise and practical best practices as critical to AI safety and effective governance.
- Encouraging proportional, risk-based governance tailored to the lifecycle stage and context of use.
- Facilitating iterative dialogue and engagement between regulators and industry to address emerging challenges.

This foundational work may be expanded into a broader report of case studies in future, offering additional depth and recommendations for regulatory consideration.



# ANNEX

## Illustrating effective AI governance

The selection of early case studies provided in this report illustrates how AI is being used and governed across different stages of the medicines lifecycle. These examples demonstrate that best practices such as independent validation, auditability, and data quality safeguards can be applied across diverse AI use cases.

A summary of AI governance across key stages of the medicines lifecycle is provided below. Rather than being a summary of the early case studies provided, it instead provides a view of where AI is being explored across the medicines lifecycle, and where EFPIA may be interested in exploring further case studies.

### Non-clinical research

AI should enhance predictive accuracy in early drug testing, significantly reducing risks of later-stage clinical failures. By reliably forecasting human physiological responses based on animal studies, AI-driven methodologies may enable safer and more efficient selection of drug candidates.

### Early clinical research

AI technologies should optimise clinical trial designs and dosing strategies, directly enhancing patient safety and reducing unnecessary exposure. AI-supported digital biomarkers may provide precise, real-time data, ensuring meaningful and accurate trial outcomes.

### Late-stage clinical trials

AI tools should streamline trial site selection and patient stratification, increasing the efficiency and accuracy of patient enrolment and treatment strategies. Real-time predictive analytics may inform patient dosing decisions, significantly improving personalised care and trial robustness. Additionally, AI-based digital endpoints should complement traditional outcome measures, providing comprehensive assessment and treatment efficacy and safety.

### Manufacturing

AI systems should enhance pharmaceutical manufacturing through real-time process control, quality monitoring, and adaptive optimisation. These applications may ensure consistent product quality, compliance with regulatory standards, and improved scalability and efficiency, directly benefiting patient safety.

### Regulatory submission

Generative AI should automate and enhance accuracy in regulatory documentation, reducing errors, streamlining processes, and accelerating regulatory review timelines. This may facilitate efficient regulatory interactions and approval processes.

### Post-approval activities

AI should significantly advance pharmacovigilance efforts, enabling early and accurate detection of adverse events. Post-market safety and efficacy studies may leverage AI-driven analytics, monitoring real-world data for continuous evaluation of medicine performance and risk identification. AI monitoring of medicine utilisation should further enhance patient adherence and treatment optimisation.

## Interview guide: Capturing best practices in AI governance across the AI lifecycle

The guide is designed for pharma executives that can provide real insights on governance, compliance, and lessons learned.

Format to help us tell a compelling story:

- Moves from business case, to governance, to impact
- Alignment with NIST AI lifecycle, while keeping it pharma-relevant
- Balance technical and governance aspects to capture best practices

EU policymakers will get a clear picture of best practices, regulatory gaps, and how to refine AI oversight.

**Objective:** Understand how pharma companies govern AI across its lifecycle, focusing on practical best practices, regulatory alignment (EMA, EU AI Act), and business benefits.

**Format:** A structured interview (primary questions with secondary questions if additional detail is needed). Follow ups will be required, both written and face-to-face.

### Plan & design: Setting the foundation for trustworthy AI

Capture how AI initiatives are planned, governed, and aligned with EU regulations & ethics from the start.

1. Can you describe the AI use case and the problem it aims to solve?
  - The fit within drug discovery, clinical trials, manufacturing, regulatory approvals, or post-market surveillance
  - The main drivers behind adopting AI for this use case
2. How do you establish governance at the planning stage?
  - Legal, regulatory, and ethical considerations taken into account (e.g., EMA guidance, EU AI Act, GDPR, transparency requirements)
  - The key decision-makers (internal and external) in ensuring AI compliance and trustworthiness
3. What steps do you take to ensure AI aligns with patient safety, fairness, and sustainability goals?
  - Consideration of potential risks to patients, clinicians, regulators, and other stakeholders
  - Trade-offs between innovation and compliance

### Data collection & processing: Ensuring high-quality, bias-free inputs

Uncover best practices in data governance, bias mitigation, and transparency.

4. What data sources were used to train the AI model?
  - Steps taken to ensure data integrity, representativeness, and compliance with EU regulations (e.g., GDPR, data anonymization)
  - Challenges in harmonising real-world data, clinical trial data, or regulatory datasets
5. How do you address bias and ensure data fairness?
  - Bias detection or mitigation techniques applied
  - Privacy-enhancing technologies, synthetic data, or federated learning to improve data security and fairness
6. How do you document and validate dataset quality?
  - Structured metadata standards or data lineage tracking tools in place
  - Handling of missing or incomplete data

## Model development & testing: Ensuring performance & reliability

Showcase how companies develop, test, and validate AI models to meet regulatory expectations.

7. How was the AI model built and trained?
  - Kind of algorithms and techniques that were used
  - Balance between performance, explainability, and regulatory requirements
8. How do you verify and validate model outputs?
  - Validation tests performed by independent teams separate from the model developers
  - Ensuring reliability, reproducibility, and generalizability of results
9. How do you balance AI explainability and performance in high-stakes decisions?
  - Using explainability techniques (e.g., SHAP, LIME) to make models interpretable
  - Feedback or concerns from regulators about model transparency

## Deployment & use: Scaling AI while managing compliance

Learn how AI is deployed in a compliant, user-friendly way while managing risks.

10. How did you ensure AI was deployment-ready in a regulatory environment?
  - Compliance tests performed before rollout
  - Addressing integration challenges within existing pharma workflows
11. How do you manage change and user adoption?
  - Training and supporting staff (e.g., medical affairs, clinical researchers, regulatory teams) in using AI-driven tools
  - Unexpected challenges in gaining user trust
12. How did AI impact efficiency, decision-making, and business outcomes?
  - Measurable benefits in productivity, cost savings, or compliance improvements
  - Benefits compare to traditional approaches

## Ongoing monitoring & risk mitigation: Continuous oversight of AI

Highlight how companies ensure AI remains reliable, unbiased, and compliant over time.

13. How do you monitor AI post-deployment?
  - Mechanisms for real-time monitoring, detecting drift, and revalidating the model
  - Unexpected risks or failures occurred, and how were they addressed
14. How do you handle regulatory audits, updates, and AI lifecycle changes?
  - Proactive updates to AI models to align with new EMA or EU AI Act guidelines
  - Documentation maintained for regulatory audits
15. What lessons have you learned from AI implementation?
  - Changing one thing about AI governance in pharma
  - Advice would for policymakers to improve AI oversight

## **Closing: Shaping AI policy & best practices**

Gather insights for EU policymakers on how AI governance should evolve.

16. What are the biggest gaps in current AI regulations for pharma?

- Conflicts or uncertainties in regulatory guidance
- Areas where more flexibility or clearer guidance is needed

17. How should policymakers support AI governance best practices?

- Policy changes or incentives that would help accelerate responsible AI adoption
- Regulatory collaboration with industry to ensure compliance while fostering innovation