


Governance and Accountability in AI-Enabled Pharmacovigilance

Dr. Sumit Verma MD, DNB - Director and President,
Operations Management, CSPV




Governance and Accountability in AI-Enabled Pharmacovigilance

- Governance and accountability are foundational to the use of artificial intelligence in pharmacovigilance. AI systems do not exist outside the pharmacovigilance system in which they are deployed. Decisions about their design, deployment, oversight, and continued use must therefore be governed with the same rigour as other safety-critical pharmacovigilance processes.
- Governance provides the framework through which decisions are made, documented, reviewed, and revised across the lifecycle of AI use.
- Accountability ensures that responsibility for those decisions remains clearly assigned, enforceable, and retained by the organisation. Together, they ensure that AI supports pharmacovigilance objectives rather than undermining patient safety or regulatory compliance.

Accountability Remains with the Organization

- The use of AI does not transfer responsibility to the technology, the algorithm, or the vendor. Accountability for pharmacovigilance outcomes remains with the organisation deploying the AI system.
- This accountability includes responsibility for:
 - Defining intended use and boundaries
 - Determining acceptable risk
 - Approving oversight models and controls
 - Deciding whether continued use remains justified

Technical activities may be delegated.
Accountability cannot.



Governance Across the AI Lifecycle

➤ Effective governance must be exercised across the full lifecycle of AI use. The nature of governance decisions varies by lifecycle phase, but responsibility and oversight must remain continuous.

1. Collection of Specifications and Requirements

Governance begins before development or configuration. At this stage, the organisation defines why AI is being used and **how** it is expected to support pharmacovigilance activities.

Key governance responsibilities include:

- Defining the intended pharmacovigilance use case and explicit exclusions
- Assessing whether AI is appropriate given the influence and consequence of error
- Defining expected outputs, limitations, and acceptable error tolerance
- Establishing the human oversight model and escalation pathways
- Defining transparency expectations for users and stakeholders

These decisions set the foundation for all subsequent governance activities.

2. Development and Change Management

During development and configuration, governance ensures that implementation choices remain aligned with approved requirements.

For systems such as large language models, behaviour is often shaped through configuration rather than model training. Prompt templates, retrieval mechanisms, constraints, and post-processing logic therefore require governance oversight.

Key controls at this stage include:

- Approval and version control of prompts and configurations
- Definition of who may make changes and under what conditions
- Documentation of how configuration choices affect system behaviour
- Formal change management for modifications that may affect outputs or risk

Changes introduced at this stage are governance events, not purely technical adjustments.

3. Pre-Deployment Sign-Off

Before deployment, governance must confirm that the system is fit for its intended use and that risks are understood and controlled.

This includes:

- Review and approval of risk assessments and mitigation measures
- Confirmation that validation and performance evaluation are appropriate to the use case
- Assessment of failure modes, including false negatives
- Approval of human oversight arrangements and user training
- Formal sign-off by accountable pharmacovigilance leadership

Deployment should only proceed once governance responsibilities at this stage have been fulfilled.

4. Post-Deployment “Hyper-Care”

Immediately following deployment or significant change, governance focus shifts to close observation of real-world behaviour.

This phase recognises that pre-deployment testing may not capture all operational risks.

Governance activities during hyper-care include:

- Enhanced monitoring of outputs and user interaction
- Rapid identification and escalation of deviations or unexpected behaviour
- Review of incidents and corrective actions
- Adjustment of oversight intensity where needed

This phase is particularly important for systems with variable behaviour, where early detection of issues is critical.

5. Routine Use and Ongoing Oversight

Once systems move into routine use, governance responsibilities continue rather than diminish.

Key activities include:

- Periodic review of system performance and impact
- Monitoring for drift, emerging failure patterns, or increased reliance
- Review of changes, updates, or reconfiguration
- Reassessment of whether oversight remains proportionate to actual influence
- Decisions to restrict, suspend, or discontinue use where appropriate

These activities should be embedded within existing pharmacovigilance quality systems.

➤ Human-in-Command as the Anchor of Governance

Human-in-command oversight underpins governance across all lifecycle phases. The organisation must retain authority to approve use, modify controls, and intervene when risks become unacceptable.

Without this authority, governance becomes symbolic rather than effective. Human-in-command ensures that accountability is not diffused across systems, vendors, or organisational boundaries.

Why LLM “Explanations” Are Not an Audit Trail

Generative language models can be prompted to produce explanations for their outputs. These explanations may appear coherent and convincing, but they do not provide direct insight into how the model arrived at a decision. They are language outputs generated after the fact and may not faithfully reflect the underlying computation.

For pharmacovigilance use, organisations should therefore prefer **evidence-anchored justifications**, such as explicit links to source text, highlighted narrative spans, or referenced documents. Where LLM-generated rationales are used, their reliability should be assessed through **faithfulness checks** before they are accepted as part of the audit trail.

➤ **Governance Considerations Specific to Large Language Models**

Large language models introduce additional governance considerations due to their variable behaviour and reliance on configuration rather than direct model training.

In such systems:

- Prompts and constraints are governed artefacts
- Changes to configuration can materially alter behaviour
- Transparency about AI involvement is essential for effective oversight
- Reliance on outputs may increase gradually and must be monitored

These considerations do not alter accountability. They clarify how governance must be exercised to remain effective.

- AI does not distribute accountability.
- It concentrates it at the points where decisions, changes, and reliance occur.
- Governance that ignores this reality creates blind spots rather than control.
- If an LLM can “explain” its output, that does not mean the explanation is an audit trail.
- Governance fails when convincing language replaces traceable evidence.

In pharmacovigilance, accountability depends on what can be verified, not what sounds reasonable.

This post continues our series on **Artificial Intelligence in Pharmacovigilance: What Safety Leaders Need to Know**

Selected References

1. Council for International Organizations of Medical Sciences (CIOMS). Artificial Intelligence in Pharmacovigilance. CIOMS Working Group XIV Report, Geneva, 2025.
2. European Medicines Agency (EMA). Reflection paper on the use of Artificial Intelligence (AI) in the medicinal product lifecycle; 2024.
3. International Council for Harmonisation (ICH) Harmonised Guideline Q9 (R1): Quality Risk Management. 2023.

Author




Dr. Sumit Verma MD DNB


Director and President,
Operations Management, CSPV
Soterius, Inc.

Dr. Sumit Verma is a medical graduate with specialization in anesthesiology and has more than 15 years of experience in the pharmaceutical industry, clinical medicine, clinical research, and pharmacovigilance. He has built teams that have consistently delivered and exceeded customer expectations across pharmacovigilance domains such as case processing, signal management, risk management, aggregate reports, and clinical safety. He has co-authored two books – one on pharmacovigilance and another on pharmacology.

Like This Blog?


 [Share on LinkedIn](#)

 [Share on Facebook](#)

 [Share on Twitter/X](#)

 [Share on Email](#)

 www.soterius.com

 connect@soterius.com

Disclaimer

Soterius logo are trademarks or registered trademarks of Soterius in all jurisdictions. Other marks may be trademarks or registered trademarks of their respective owners. The information you see, hear or read on the pages within this presentation, as well as the presentation's form and substance, are subject to copyright protection. In no event, may you use, distribute, copy, reproduce, modify, distort, or transmit the information or any of its elements, such as text, images or concepts, without the prior written permission of Soterius. No license or right pertaining to any of these trademarks shall be granted without the written permission of Soterius (and any of its global offices and/or affiliates). Soterius reserves the right to legally enforce any infringement of its intellectual property, copyright and trademark rights.

The content presented is intended for general informational purposes only. It does not constitute legal, regulatory, medical, or other professional advice, and should not be relied upon as a substitute for such advice. The information is not specific to the circumstances or requirements of any particular organization or use case. While care has been taken in preparing this material, Soterius, including its authors and affiliated contributors, makes no representations or warranties regarding the completeness, accuracy, reliability, or applicability of the content. Nothing in this content is intended to provide medical advice or clinical recommendations. Organizations and individuals should seek independent professional advice before acting on any information presented.