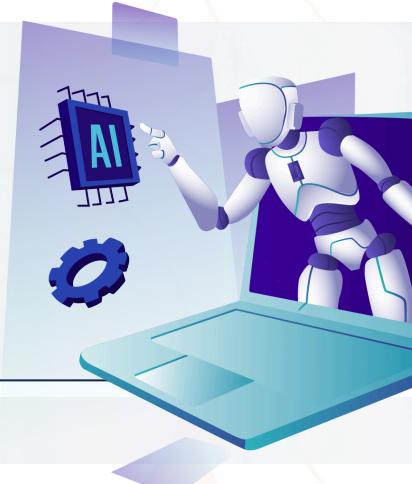


# Deploying an Inspection Ready AI System in PV

**Artificial intelligence in pharmacovigilance should not be assessed by how advanced the technology is. It should be assessed by what happens when the technology is wrong.**

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



## Executive Summary

- Artificial intelligence is increasingly being applied across pharmacovigilance activities, particularly in case processing, triage, and assessment. While these technologies are often introduced to improve efficiency and scalability, their use raises important questions about risk, oversight, and accountability.
- In a GxP-regulated environment, AI systems used in pharmacovigilance must meet regulatory expectations for risk management, oversight, and accountability.
- It is important to develop a risk-based perspective on the use of artificial intelligence in pharmacovigilance. AI systems should not be evaluated as standalone productivity tools, but as components of the pharmacovigilance system whose outputs can materially influence regulatory decisions and patient safety outcomes.
- Risk in AI-enabled pharmacovigilance is defined not by algorithmic complexity, but by two factors:
  - Degree to which AI outputs influence decisions &
  - Consequences of incorrect outputs.Even technically simple systems may represent high risk when their outputs affect reporting obligations, safety assessments, or downstream signal detection activities.
- We should understand the known and foreseeable failure modes, including false negatives, automation bias, and the compounding effect of downstream reuse of AI outputs. It is important to know why human oversight must be deliberate and proportionate to risk, and why human-in-the-loop models alone are insufficient without clear accountability and safeguards against over-reliance.
- Finally, emphasis should be placed on the role of safety leadership. Classification of AI risk, design of oversight models, and continuous monitoring are leadership responsibilities that cannot be delegated to technology providers or implementation teams. As AI systems become more influential within pharmacovigilance workflows, accountability becomes more concentrated, not diluted.  
Organizations should consider applying a structured, inspection-ready approach to AI in pharmacovigilance, grounded in risk awareness, transparency, and patient safety.

## Why AI in Pharmacovigilance Must Be Treated as a Risk-Based System

- Artificial intelligence in pharmacovigilance should not be assessed by how advanced the technology is. It should be assessed by what happens when the technology is wrong.
- This is a deliberate position. AI used in pharmacovigilance must be evaluated primarily through the lens of risk, not innovation.
- Risk is driven by two factors:
  1. Consequence of an incorrect output &
  2. Degree to which the AI system influences pharmacovigilance decisions.Sophistication or novelty does not reduce risk. Influence and impact define it.
- This framing matters for Heads of Safety. An AI system used for case processing, triage, or assessment is not evaluated as a productivity enhancement. It is evaluated as part of the pharmacovigilance system itself, with direct implications for patient safety, regulatory compliance, and organisational accountability.
- Risk assessment is therefore not a one-time exercise performed at deployment. It must be applied throughout the lifecycle of AI use.
- This is why AI in pharmacovigilance must be treated as a risk-based system, not a tool.

## How Risk Is Defined in AI-Enabled Pharmacovigilance

- Risk in AI-enabled pharmacovigilance is not determined by how complex the technology appears to be. It is determined by how much influence the AI system has on decisions and the consequence of an incorrect decision.
- This distinction is critical. An AI system performing a technically simple task can represent high risk if its output directly affects regulatory reporting, medical judgement, or patient safety. Conversely, a technically complex system may present lower risk if its outputs are clearly limited, well understood, and subject to effective human control.
- Risk increases as AI systems move closer to stand-alone operation. When outputs are consumed without meaningful challenge, or when they materially shape decisions without sufficient human intervention, the potential impact of error increases. In pharmacovigilance, even small errors can alter the understanding of a medicine's benefit-risk profile.

- Risk is also contextual. The same AI capability can fall into different risk categories depending on how and where it is used. An extraction model supporting case intake may be lower risk when outputs are fully reviewed, but higher risk when those outputs are reused downstream to drive assessments or reporting decisions.
- For Heads of Safety, this means risk classification is not a technical exercise delegated to vendors or data scientists. It is a business and safety decision that must consider real workflow influence, error detection mechanisms, and actual reliance on the system in practice. This understanding underpins expectations for oversight, monitoring, documentation, and governance.

## Explicit Risks That Must Be Anticipated

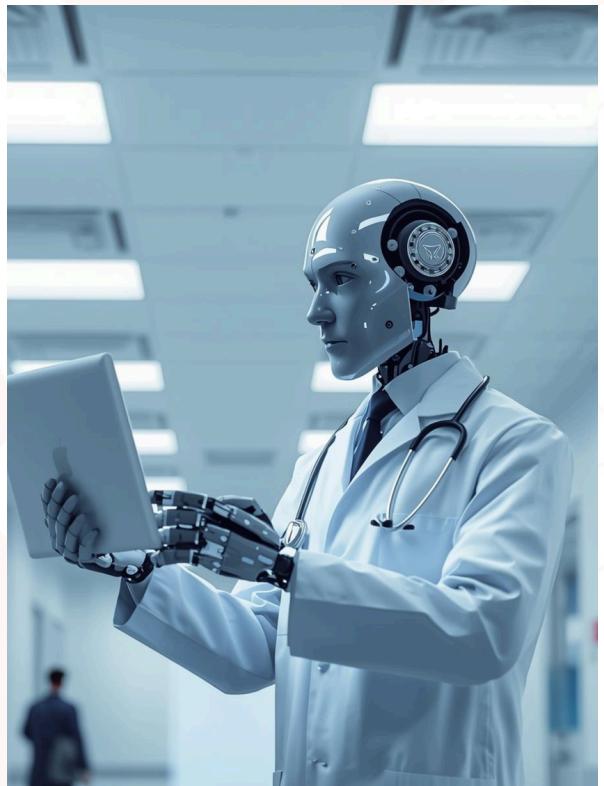
- The risks associated with AI in pharmacovigilance are not theoretical. Specific failure modes are well understood and particularly relevant to safety-critical activities.  
One of the most significant risks is false negatives. AI systems may fail to identify relevant adverse events, seriousness criteria, or safety signals, even when overall performance metrics appear acceptable. In pharmacovigilance, these failures are not benign. Missed information can delay reporting, distort safety assessments, and undermine patient protection.
- Automation bias represents another important risk. When AI systems are perceived as reliable, human reviewers may place undue trust in their outputs. This can lead to reduced vigilance, insufficient challenge, and a gradual erosion of critical review. Human review alone does not eliminate this risk. Without deliberate safeguards, human involvement can become procedural rather than protective.
- A further risk arises from downstream reuse of AI outputs. Information generated early in the workflow is often reused in later steps such as assessment, aggregation, or signal detection. As outputs travel further through the pharmacovigilance process, their influence increases and errors can compound, even if the original use case appeared limited.
- These risks increase whenever AI outputs influence regulatory or clinical judgement. High average accuracy or good intentions are not sufficient safeguards. Known failure modes must be anticipated and actively mitigated.
- For Heads of Safety, the implication is clear. These risks are foreseeable. Failure to address them cannot be justified as unexpected system behaviour.

## What This Means for Case Processing and PV Use Cases

- Many AI use cases in pharmacovigilance are introduced as operational support. In practice, their influence often extends far beyond clerical assistance. Case processing is a clear example. AI systems used to extract medical concepts, identify seriousness criteria, or prioritise cases directly influence regulatory timelines and downstream safety evaluation. Even when described as assistive, these outputs frequently shape what is reviewed, what is escalated, and what is ultimately reported.
- Errors in these contexts may be infrequent but consequential. A missed hospitalization, a misinterpreted narrative, or an overlooked adverse event may not significantly affect aggregate metrics, but it can materially affect patient protection and compliance.
- Another compounding factor is downstream reuse. Risk accumulates not because the algorithm changes, but because its outputs travel further through the pharmacovigilance process.
- As a result, many AI-supported case processing activities sit higher on the risk spectrum than initially assumed. Classification cannot be based on convenience. It must be based on influence and consequence.

## Human-in-the-Loop and Proportionate Oversight

- Human oversight is not symbolic review. It is an active control designed to mitigate known risks.
- One important oversight model is human-in-the-loop. In this approach, AI outputs are reviewed and either accepted, modified, or rejected by a qualified human before they influence decisions.
- Human review alone, however, does not guarantee accuracy. Automation bias can undermine oversight when reviewers unconsciously defer to AI outputs. Without deliberate safeguards, human-in-the-loop can degrade into human-on-paper. Oversight must therefore be proportionate to risk. As AI influence increases, so must the clarity of reviewer responsibility, the depth of review, and the supporting quality controls. Oversight models must account for human behaviour, not just system performance.



## What This Means for Safety Leadership and Accountability

- Risk classification of AI use in pharmacovigilance is not a technical decision. It is a leadership responsibility.
- Safety leadership must understand how AI systems are actually used, how outputs influence decisions, and where errors could have impact. Vendor descriptions, average accuracy, or intended use statements are insufficient if they do not reflect real operational behaviour.
- Risk assessment is not static. AI systems evolve, workflows change, and reliance can increase over time. Continuous review, monitoring, and escalation mechanisms are therefore essential.
- This is why expectations around human oversight, performance monitoring, and governance exist. They are not generic best practices. They are responses to known and foreseeable risks.
- AI does not dilute accountability. It concentrates it. Artificial intelligence will increasingly shape pharmacovigilance activities. The risks are not hypothetical, and ignoring them is not an option.

**In subsequent posts, we will explore how these risks can be mitigated in practice, including the role of human-in-the-loop oversight, monitoring, and governance.**

## 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](http://www.soterius.com)[connect@soterius.com](mailto: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.**