WHITE PAPER

Artificial Intelligence in Pharmacovigilance: The Present and the Future

Improving efficiency, consistency, and compliance across the pharmacovigilance lifecycle, with human oversight at the core.

Dobrochna Dolicka · **Christopher Henry** · **Laurent Vidal** · **Alexis Pinçon**

Global Safety Writing & Medical Services, Pharmacovigilance, UBC

Introduction

Artificial intelligence (AI) is reshaping pharmacovigilance (PV), by improving efficiency, consistency, and compliance across multiple domains. AI technologies are being integrated into many PV activities, including:

- case intake
- triage
- causality assessment
- periodic safety reporting
- REMS annual report preparation
- literature review
- regulatory updates

As adoption grows, organizations must balance automation versus AI and determine when to build or buy solutions based on operational and compliance needs. Looking ahead, PV is moving toward a “centaur model,” combining human expertise with AI-driven insight to enable more proactive, predictive safety monitoring while maintaining human oversight.

This piece outlines current applications, key limitations, and future opportunities for AI in PV, with a focus on practical, compliant implementation.

> *AI is becoming essential and not optional, in modern PV operations.*

In such a highly regulated environment where accuracy is critical, organizations must understand both the value and the limitations of AI. This article discusses real-world and emerging AI use cases across pharmacovigilance (PV) and introduces the “centaur” model, where human expertise and AI work collaboratively to improve speed, quality, and decision making.

Use Cases: Periodic Safety and REMS Reporting

Periodic Safety Reports

Periodic safety reporting (DSUR, PSUR) is inherently data-intensive, requiring the synthesis of large, multi-source datasets under tight timelines. AI is increasingly used to automate data aggregation, identify safety trends, and generate structured draft content, but also identify inconsistencies within the report.

Rather than replacing scientific input, these tools can shift effort toward interpretation and validation, enabling teams to focus on benefit-risk evaluation and regulatory positioning.

While AI can support drafting and data reconciliation, final accountability, interpretation, and stakeholder communication remain firmly human-led.

REMS assessment reports

Risk Evaluation and Mitigation Strategies (REMS) assessment reports involve recurring updates, structured datasets, and standardized outputs, making it suitable for AI augmentation.

AI could automatically integrate updated tables into reports, synchronize associated narrative text, and identify changes from prior submissions. In more advanced cases, it can generate sections directly from structured inputs, significantly reducing manual effort.

The result? Improved efficiency and reduced risk of inconsistency, particularly in large or complex programs.

Use Cases: Case Processing

AI is having a transformational impact on case processing, where volume, variability, and regulatory timelines are intersecting.

Case narrative

Narrative writing, traditionally time intensive, can now be accelerated through AI generated drafts, based on structured safety data. These outputs provide a consistent baseline, allowing for medical writers to focus on refinement, clinical nuances and complex cases.

Human review remains important and essential, particularly to ensure accuracy, context and regulatory compliance.

Triage

In pharmacovigilance, case triage is a crucial initial step that involves assessing incoming ICSRs to determine their validity, seriousness, expectedness, and reportability. AI can enable rapid classification and prioritization based on predefined regulatory criteria, reducing delays and improving overall consistency.

In further practice, AI could enhance throughput while flagging incomplete or ambiguous cases, ensuring that human reviewers are focused on high value interventions, improving urgent cases needing submission.

Causality assessment

Causality assessment remains one of the most nuanced areas in PV. AI could support the process by analyzing historical patterns, identifying associations and generating probability-based insights with supporting rationale. These outputs serve strictly as decision support and final determination requires clinical judgement, contextual awareness, and accountability that AI simply cannot replicate.

Use Cases: Case Processing (continued)

Data entry from source documents

Advances in OCR and natural language processing have made it possible to extract structured data from diverse and unstructured source materials, including handwritten or multilingual inputs.

This enables a shift from manual data entry to AI-driven population of safety databases, with the added capability to identify missing information and prompt follow-up.

The impact is both operational and qualitative and can aid in reducing administrative burden while improving data completeness.

PT coding

AI can not only enhances MedDRA coding by understanding ambiguous verbatim terms, but also by resolving inconsistencies, and suggests appropriate Preferred Terms based on context.

This can improve both speed and consistency, particularly at scale, while retaining human oversight for ambiguous or complex cases.

Literature review

Across both global and local contexts, AI can significantly improve literature review by automating identification, assessment, and extraction of safety-relevant information or ICSRs. These tools can process large volumes of multilingual content, identify potential signals, and even generate new cases where appropriate.

This shifts the role of PV teams from manual screening to higher-value signal detection and analysis.

Use Cases: Regulatory Intelligence and Operational Support

Safety Regulatory Intelligence

Safety regulatory intelligence is essential in a heavily regulated space such as pharmacovigilance. For small pharma companies, this can be relatively straightforward but when the number of countries where clinical trials are ongoing or where the drug is approved reaches a certain number, regulatory monitoring becomes increasingly challenging due to country specificities. Monitoring regulatory requirements across jurisdictions is resource-intensive and constantly evolving. AI can support this function by tracking updates, analyzing regulatory changes, and maintaining current intelligence across regions.

In parallel, AI can streamline submission preparation by generating standardized documents and validating completeness against regional requirements.

Together, these capabilities improve compliance, reduce manual effort, and enable faster response to regulatory change.

Submission package preparation

AI technologies can easily be integrated into the preparation of report submission packages in pharmacovigilance, streamlining the generation of cover letters and submission-ready documents. AI could generate customized cover letters by extracting key data points such as product names, reporting periods, regulatory authority details and safety findings, from periodic safety reports. These systems could adapt templates based on specific health authority requirements and ensure consistency in language, tone, and structure. Additionally, AI tools could assist in compiling annexes, formatting the package according to electronic submission standards (e.g., Electronic Common Technical Document [eCTD]), and validating the completeness of required documents before submission. This would not only reduce human error and save significant time but would also improve compliance with region-specific submission guidelines.

General PV-oriented support

Beyond core safety workflows, AI is increasingly embedded in day-to-day operational support.

It enhances project management by summarizing meetings, structuring action items, and tracking follow-ups, enabling teams to operate more efficiently. Similarly, AI-powered translation tools improve speed and consistency across multilingual workflows, though human validation remains essential in regulated environments.

AI can also enables the transformation of SOP libraries into interactive, searchable knowledge systems, improving access to procedural guidance and supporting onboarding, particularly handy when product specificity complexifies the picture.

In analytics, AI introduces pattern recognition and adaptability, complementing traditional automation and allowing teams to derive deeper insights from complex safety data to support their signal management or aggregate reporting activities, among others.

Limitations: IT Security

Despite clear benefits, AI must be implemented with careful consideration of risk, compliance, and operational constraints. AI systems introduce unique IT security challenges because of their large scale and complex functioning.

Cloud-based solutions

Cloud-based solutions enable faster data access and collaborative work environments. However, they also introduce a new set of security and compliance challenges, mainly due to unharmonized data protection regulations in different jurisdictions, as well as associated risks of service disruptions and data breaches.

Mitigation strategies include selecting cloud providers with strong compliance credentials, implementing end-to-end encryption, and using private or hybrid cloud architectures.

Local infrastructure

On the other hand, local infrastructure offers an appealing level of control over pharmacovigilance data and workflows, particularly in highly regulated environments. Hosting sensitive safety data internally reduces reliance on external vendors and allows for customized security configurations aligned with internal policies.

However, this control comes with limitations in scalability, maintenance, and resource allocation, as well as internal security issues.

Hybrid approach

Balancing autonomy with efficiency often requires a hybrid approach, leveraging the robustness of local solutions for critical data, while using controlled cloud extensions to enable AI-driven innovation.

Limitations: Regulatory and Privacy Considerations

Data privacy remains a central concern in PV, particularly given AI's ability to be content generative from processing large volumes of structured and unstructured data. Compliance with frameworks such as HIPAA and GDPR requires rigorous safeguards, including de-identification or anonymization, and clear governance structures.

Human oversight is essential, particularly in scenarios involving automated decision-making or patient-level data.

Limitations: Quality Assurance and Validation

Over the years, comprehensive standards have been developed, however, the appearance of AI introduces new challenges to ensure a robust validation of the system scope of use.

Risk based approach

Given the complexity and evolving nature of AI models, particularly those based on machine learning, traditional validation frameworks can prove inadequate or unsustainable. A risk-based approach has emerged as a pragmatic alternative - focusing validation efforts on high-impact functionalities while allowing for adaptive learning in lower-risk areas.

Depending on the level of risk associated with the AI system use, more human oversight can be added along with a strong transparence to what the system is used for and its possible limitations, including risk mitigation solutions when needed.

While this risk-based method provides flexibility and reduces validation burdens, it introduces its own set of limitations. Determining risk thresholds can be subjective and may vary across stakeholders, leading to inconsistent implementation. Moreover, regulators may challenge the justification of such an approach, especially in systems affecting safety case processing or signal detection. Different groups are trying to harmonize this approach such as FDA [2], EMA [5], ICH [3], and the Council for International Organizations of Medical Sciences (CIOMS) work practice [1].

Proper documentation, audit trails, and layered testing are essential to ensure traceability. However, ongoing alignment with evolving regulatory expectations remains critical to maintain trust in AI-supported PV systems.

Human in the loop

AI in pharmacovigilance is rarely deployed in full autonomy. Most validated use cases operate within a human-in-the-loop or human-on-the-loop framework, where expert oversight is maintained either actively during the AI task (human-in-the-loop) or passively via post-review (human-on-the-loop). These hybrid models are designed to combine the efficiency of AI systems with the accountability and contextual judgment of human reviewers.

However, their effectiveness hinges on clearly defined roles and seamless integration. If human reviewers are overwhelmed with volume or unclear about when intervention is necessary, the system risks either underperformance or overdependence on manual checks - both of which undermine efficiency.

Limitations: Quality Assurance and Validation (continued)

Testing

Testing is a foundational element in AI validation, but its role shifts significantly when we move from deterministic scripts to probabilistic models.

Unlike scripted logic, where outputs are predictable and testable with near-100% accuracy, AI models introduce variability by design. Their behavior can fluctuate based on input nuances, data drift, or underlying probabilistic logic. This makes it difficult to define fixed pass/fail criteria, especially in safety-critical domains like pharmacovigilance.

Traditional validation often relies on static datasets that may not capture the real-world diversity or rare case patterns the model will encounter. As such, validating the full AI-enabled workflow - rather than just the model - is crucial. Furthermore, models are constantly evolving; when ChatGPT 5 was released, it was proved to be far superior to ChatGPT 3.5, it is therefore not possible to simply upgrade the models used without additional verifications. This calls for dynamic quality assurance (QA) strategies that include ongoing performance monitoring, anomaly detection, and rapid escalation protocols to ensure that the best model is always in use.

In this paradigm, validation is not a one-time gate but a continuous assurance process acknowledging that 100% accuracy is neither expected nor realistic.

Discussion

When to consider AI versus automation?

The choice between AI and automation is not binary.

Automation

Automation remains highly effective for structured, predictable, rule-based tasks.

- "Fill in the blank"
- Recurring same tasks
- Analyses based on structed raw data
- Complex calculations

AI

AI excels in context-driven, variable, and interpretive scenarios.

- "Creative"
- Organizing content from scratch
- Out of the box analyses

Figure 1. Automation versus AI

In practice, the most effective PV operations integrate both, leveraging automation for efficiency and AI for adaptability.

What to develop and what to outsource?

Organizations must balance cost, scalability, and specificity when selecting AI solutions. Commercial tools offer speed, reliability, and ease of implementation, while in-house solutions provide greater flexibility and long-term strategic value, particularly in niche or highly specialized use cases.

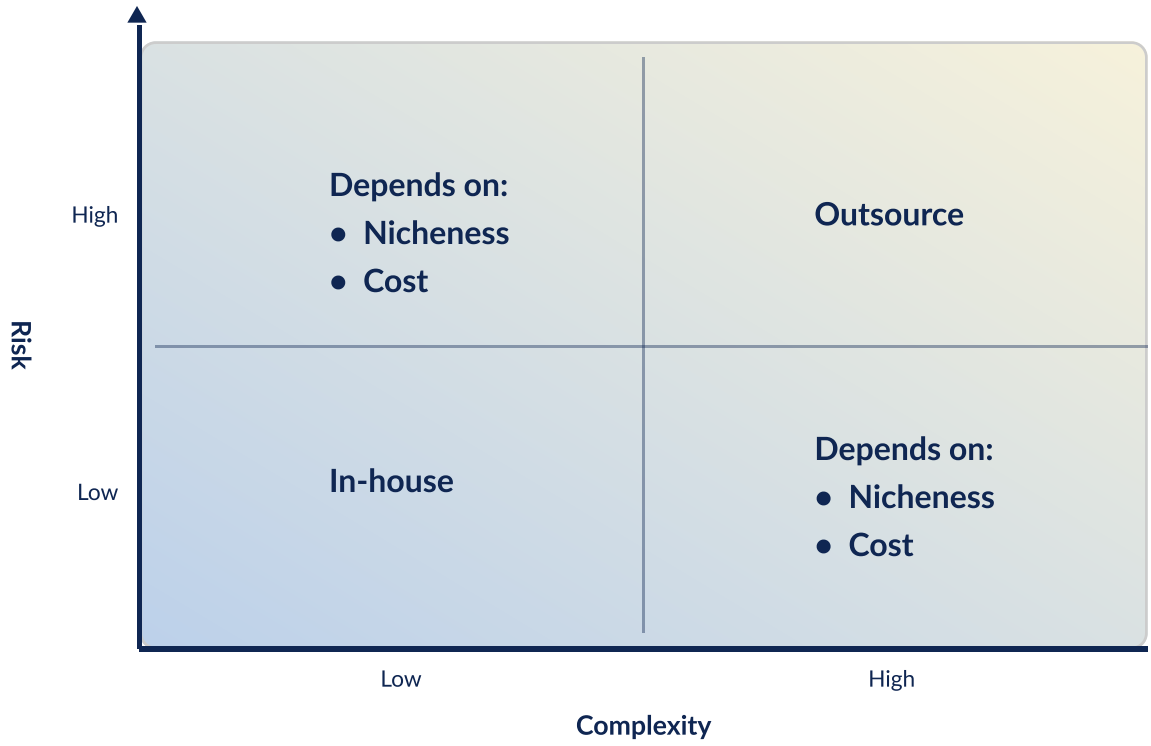

Figure 2. What to develop and what to outsource depends on the risk and complexity of the tool, as well as its nicheness and total cost.

What should not be delegated to AI?

At present, it is not yet feasible to delegate all decision-making tasks to artificial intelligence. Certain responsibilities, particularly those involving clinical judgment and regulatory accountability, must remain human-led.

While AI can highlight relevant insights or patterns, it cannot fully interpret nuance or assume responsibility for safety-critical decisions. Overreliance also introduces risk, including cognitive bias.

Maintaining strong human expertise is essential to safe and effective AI adoption [1].

The Future of PV: The Centaur Model

The future of pharmacovigilance lies in intentional collaboration between humans and AI.

The centaur model

AI	Human experts
analysis, interpretation, and insight generation	judgment, oversight, and accountability

This “centaur model” [4] represents a shift toward more proactive, predictive, and scalable PV systems, while preserving trust and patient focus.

> *The goal is not replacement, but augmentation. Redefining how expertise is applied in an AI-enabled environment.*

Ultimately, pharmacovigilance will become more proactive and predictive, powered by real-time data streams, continuous learning systems, and human oversight that ensures accountability, trust, and patient safety. These large-scale changes will also require careful change management planning to ensure a smooth transition and effective adoption of AI across teams. For a deeper dive on this aspect, have a look at our thought process, in “Redefining PV Skillsets in an AI-Driven World: Training the Next Generation of Safety Scientists”.

References

[1] Artificial intelligence in pharmacovigilance. CIOMS Working Group report. (n.d.). Retrieved from https://cioms.ch/wp-content/uploads/2022/05/CIOMS-WG-XIV_Draft-report-for-Public-Consultation_1May2025.pdf

[2] Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products. (n.d.). Retrieved from https://www.fda.gov/media/184830/download

[3] ICH M15 guideline on general principles for model-informed drug development - Step 2b - Scientific guideline. (n.d.). Retrieved from https://www.ema.europa.eu/en/ich-m15-guideline-general-principles-model-informed-drug-development-step-2b-scientific-guideline

[4] Mollick, E. (2024). Co-intelligence. Penguin.

[5] Reflection paper on the use of artificial intelligence in the lifecycle of medicines. (n.d.). Retrieved from https://www.ema.europa.eu/en/news/reflection-paper-use-artificial-intelligence-lifecycle-medicines

From Safety to Evidence to Access

Global Safety Writing & Medical Services, Pharmacovigilance · UBC