IntuitionLabs

INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

AI Governance in Pharmacovigilance Signal Detection

> Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2026-04-05 · ai governance · pharmacovigilance · signal detection · cioms xiv · transcelerate · drug safety monitoring · machine learning · regulatory frameworks

Original article: https://intuitionlabs.ai/articles/ai-governance-pharmacovigilance-signal-detection



In this article

1. Executive Summary
2. Introduction and Background
3. TransCelerate's Initiatives in Pharmacovigilance
4. TransCelerate Overview
5. Intelligent Automation in PV
6. The Role of AI in Signal Detection – Technology Perspective
7. Traditional Signal Detection vs AI-Enabled Detection
8. Evidence of AI Effectiveness
9. Implementing AI Governance in Signal Detection
10. Core Governance Principles
11. Organizational Governance and Oversight
12. AI Model Lifecycle and Validation
13. Data Governance and Standards
14. Organizational Action Items
15. Case Study: Pilot Implementation of AI Triage
16. Regulatory Frameworks and Standards
17. Discussion and Future Directions
18. Conclusion

Executive Summary

The advent of artificial intelligence (AI) and machine learning (ML) has ushered in a transformative era for pharmacovigilance (PV), particularly in the realm of safety **signal detection**. Global pharmaceutical consortia and regulatory bodies alike recognize the potential of AI to enhance drug safety monitoring, but also the need for robust governance to manage risks. In December 2025 the Council for International Organizations of Medical Sciences (CIOMS) published its *Working Group XIV* report on **Artificial Intelligence in Pharmacovigilance** (CIOMS). This landmark report establishes a consensus "safety net" of principles—ranging from risk-based oversight and human participation through transparency, data quality, and accountability—to guide the responsible deployment of AI in PV (pvpharm.com) (sakaradigital.com). Simultaneously, TransCelerate BioPharma, an industry coalition, has produced extensive resources (surveys, case studies, frameworks) on emerging digital tools in PV (Kassekert et al.) (www.transceleratebiopharmainc.com).

These parallel initiatives underscore a convergence: achieving the promise of AI (greater speed, coverage and insight in signal detection) requires equally rigorous governance structures. AI-augmented PV can process immense volumes of Individual Case Safety Reports (ICSRs), integrate diverse data sources (electronic health records, literature, social media) and potentially improve surveillance efficiency (FDA) (sakaradigital.com) (pvpharm.com). For efforts to cut manual processing time, CIOMS Appendix 3, use case A, reports a proof of concept for LLM-assisted case intake with an estimated 39% efficiency gain and approximately 20 minutes saved per case (CIOMS final report). However, these potent tools also raise concerns: algorithmic bias, data

privacy, “black-box” uncertainty, and the risk of undetected errors with patient-safety implications. Thus, both CIOMS and TransCelerate urge that **governance be built in** – through risk-based strategies, validation protocols, continuous monitoring, and clear accountability – to ensure that AI serves **as an enabler of better signal detection without compromising patient safety** (pvpharm.com) (sakaradigital.com).

This report provides a comprehensive analysis of the current state and future direction of AI in PV signal detection, focusing on governance frameworks. We begin with background on pharmacovigilance and historic motivations for AI, then review TransCelerate’s and CIOMS’s roles in shaping best practices. We examine how AI/ML techniques (including NLP and large language models) are applied to signal detection and the evidence of their performance (link.springer.com) (sakaradigital.com). A central section dissects **AI governance**: key principles (risk-proportionate oversight, explainability, data ethics) and practical models (validation strategies, oversight committees, phased implementation) from industry and regulatory sources (pvpharm.com) (sakaradigital.com). Case studies illustrate how organizations pilot and scale AI-based signal processes under these guidelines (e.g., parallel human review during initial triage) (sakaradigital.com). Lastly, we discuss implications: regulatory trends (FDA’s Emerging Drug Safety program and the joint FDA–EMA principles for good AI practice across the medicines lifecycle, published in January 2026 (EMA–FDA announcement)), workforce readiness, and emerging innovations like federated learning and predictive safety modeling (sakaradigital.com) (pvpharm.com).

In summary, **translating AI’s promise into patient benefit mandates disciplined governance**. The guidance of CIOMS WG XIV and TransCelerate, backed by academic and industry research, provides a practical roadmap: embed human expertise at critical junctions, demand transparency and continuous validation, and align AI use with fundamental PV objectives. With these frameworks, AI-enabled signal detection can evolve PV from a reactive process to a proactive, data-driven guardian of drug safety.

Introduction and Background

Pharmacovigilance (PV) is the discipline of monitoring the safety of medical products post-launch, primarily through the collection and analysis of adverse event reports. At its core is **signal detection** – the process of identifying new or unexpected adverse drug reactions (ADRs) from routine data sources (link.springer.com). Conventionally, sponsors and regulators rely on statistical disproportionality methods (e.g. Proportional Reporting Ratio, Reporting Odds Ratio) applied to individual case safety report (ICSR) databases, as defined in ICH guidances and CIOMS frameworks (link.springer.com). Feature exploration and alternative machine-learning algorithms may strengthen signal detection (Dauner et al.). Meanwhile, the volume and heterogeneity of PV data are skyrocketing. FDA’s Office of Surveillance and Epidemiology reported receiving **2,160,304 FAERS reports in 2023** (2023 annual report); this is a dated report count for FAERS, not a count of unique patients or a total across all FDA reporting systems. New unstructured sources – electronic health record (EHR) data, medical literature, social media posts – offer vast safety information but evade classical methods without intense manual curation.

In parallel, **artificial intelligence** (AI) and **machine learning** (ML) have matured dramatically. Modern ML models (decision trees, ensemble learners, neural networks) and NLP techniques can process large, complex datasets and adapt autonomously. In PV, this means algorithms can sift through millions of reports, extract

meaning from narratives, and flag subtle patterns that humans or simple statistics might miss (sakaradigital.com). These capabilities promise to shift PV from "warm-start" retrospective analysis to **cold-start predictive surveillance**, potentially identifying risks even in early clinical stages (pvpharm.com).

TransCelerate BioPharma Inc., a non-profit consortium of major pharmaceutical companies, has been proactive in harnessing AI for PV. Its *Intelligent Automation Opportunities in PV* initiative (completed in 2025) surveyed member companies and developed practical tools on automation (including AI, robotic process automation, NLP) across PV processes (www.transceleratebiopharmainc.com) (Kassekert et al.). TransCelerate's research, published in *Drug Safety* (2022), documented a clear industry trend: life sciences organizations are rapidly moving from planning to piloting and deploying intelligent automation in case processing (Kassekert et al.). These surveys found that companies are using multiple AI techniques in tandem (e.g. ML for coding and RPA for data entry) and highlighted top challenges – chiefly the scarcity of high-quality training data and a lack of harmonized regulatory guidance for AI validation (Kassekert et al.). TransCelerate has since produced guidance assets (e.g. technology matrices, case studies on signal management automation (www.transceleratebiopharmainc.com)) to help companies navigate these challenges.

Meanwhile, CIOMS established **Working Group XIV (WG XIV)** to develop principles and guidance for AI in PV (CIOMS). This interdisciplinary group (regulators, industry, academia, patient reps) met through 2025, culminating in a public-consultation draft (May 2025) and final report (Dec 2025) on *Artificial Intelligence in Pharmacovigilance*. The CIOMS report explicitly frames AI in PV as a "safety net" challenge: it offers transformative tools but demands a commensurate ethical/operational framework. In its summary, CIOMS WG XIV introduced seven guiding principles – risk-based oversight, human oversight, validity, transparency, privacy, fairness, and governance – all designed to ensure AI enhances rather than undermines drug safety (pvpharm.com) (CIOMS guiding principles). These principles underscore that **patient safety primacy** must not be compromised by pursuit of efficiency gains (sakaradigital.com) (pvpharm.com).

Scope of Report: This document examines the intersection of TransCelerate's industry initiatives and the CIOMS WG XIV framework in the context of PV signal detection (especially by AI). We analyze current technologies and evidence in AI-enabled signal detection, then focus on **AI governance** – the policies, processes, and oversight needed to implement AI in PV safely. The goal is to present a deep, evidence-based synthesis drawing on academic studies, industry reports, and regulatory publications. We consider multiple perspectives (industry, regulatory, ethical), include quantitative data where available, and discuss real or conceptual case studies. Ultimately, we aim to elucidate how AI can be responsibly governed in PV — turning innovative models into reliable, compliant signal systems.

TransCelerate's Initiatives in Pharmacovigilance

TransCelerate Overview

TransCelerate BioPharma Inc. is a global non-profit consortium of major pharmaceutical companies founded in 2012. Its mission is collaborative innovation to accelerate the delivery of new medicines. Members share resources and best practices on clinical and regulatory science. In pharmacovigilance, TransCelerate has formed working groups to address common challenges. Notably, the **Interpretation of Pharmacovigilance Guidances & Regulations (IGR PV)** initiative produces guides to harmonize diverse global PV rules (www.transceleratebiopharmainc.com). Another key initiative has been **Intelligent Automation in PV**, aiming to harness emerging technologies (AI/ML, robotics, NLP) to improve efficiency and quality in safety reporting (www.transceleratebiopharmainc.com).

TransCelerate's value lies in combining industry expertise into consensus-driven outputs. This collective approach reduces duplication of effort and provides regulators with unified industry positions. For example, TransCelerate's IGR PV published manuscripts clarifying contentious safety reporting topics (e.g. definition of anticipated SAEs, reference safety information) (www.transceleratebiopharmainc.com). Similarly, its AI/automation work reports results from its 2019, 2020, and 2021 member-company surveys on intelligent automation in pharmacovigilance (Kassekert et al.) (www.transceleratebiopharmainc.com).

Intelligent Automation in PV

The *Intelligent Automation Opportunities in Pharmacovigilance* initiative (completed April 2025) focused on identifying how "intelligent automation" technologies can streamline PV processes (www.transceleratebiopharmainc.com). The initiative performed an **impact assessment** and worked with global health authorities to vet risks (www.transceleratebiopharmainc.com). Its deliverables include:

- **Surveys of Company Adoption (2019–2021):** As reported in *Drug Safety* (2022), TransCelerate surveyed member companies on their use of automation in PV. The findings showed rapid adoption: firms have moved from planning to pilot to production for many rule-based automations, and are increasingly applying ML/AI for tasks such as event coding and duplicate detection (Kassekert et al.).
- **Technology Matrix:** An interactive catalog classifying various AI/automation tools and their readiness for PV tasks (available on TransCelerate's website).
- **Case Study Themes:** For example, an "AI-based Validation Case Study Themes" asset describes scenarios of validating ML models in PV. (www.transceleratebiopharmainc.com).
- **Signal Management Survey:** A report "Intelligent Automation Opportunities for Signal Detection" (November 2023) summarized industry needs and experiences in automating signal management processes (www.transceleratebiopharmainc.com).
- **Implementation Tools:** Tools like the Interactive ICSR and Automation Technologies Tool (IATT) help companies match processes to suitable technologies (www.transceleratebiopharmainc.com).

- **Published Literature:** TransCelerate members have published springer articles drawing on these efforts (e.g. Kassekert *et al.*, 2022) (Kassekert et al.).

From these resources, some key industry perspectives emerge:

- **Broad Interest:** Surveyed member companies remain highly interested in AI/ML technologies that can deliver human-like interpretation of data and support decision-making in case processing (Kassekert et al.).
- **Combined Technologies:** Solutions often integrate multiple tools (e.g. ML + robotic process automation) for the same task (Kassekert et al.).
- **Primary Use Cases:** Initial focus areas include duplicate report detection, automated MedDRA coding of events, and extracting data from narrative case descriptions (link.springer.com) (pvpharm.com).
- **Challenges:** The two biggest hurdles are *training data* (acquiring sufficient, high-quality annotated data for ML) and *regulatory uncertainty*. Companies reported difficulty in obtaining representative datasets to train and validate models (Kassekert et al.). They also expressed a need for clearer regulatory guidance to know how to qualify their AI systems (Kassekert et al.).

TransCelerate's outputs are therefore geared toward addressing these hurdles: by educating PV teams on technology capabilities, benchmarking which steps can be automated, and providing draft protocols for model validation. For example, TransCelerate's repository includes an outline for *validating intelligent automation systems* consistent with quality frameworks .

Overall, TransCelerate plays a dual role: (1) **Innovator & Educator**, by pooling lessons from member pilots to advance PV technology (as evidenced in conference talks and publications), and (2) **Facilitator of Harmonization**, by engaging regulators (as it has done for other guidances) to ensure new tools fit into existing safety frameworks. In the AI era, such industry-wide collaboration helps translate the fast-moving AI frontier into robust industry standards.

The Role of AI in Signal Detection – Technology Perspective

AI in pharmacovigilance has most tangibly impacted **signal detection**, where complex data challenges align with AI strengths. This section delves into how AI/ML techniques augment traditional methods, with evidence from recent studies and practical reports.

Traditional Signal Detection vs AI-Enabled Detection

Traditionally, PV signal detection is performed by periodically analyzing the incidence of drug-event combinations in ICSR databases. Methods like the Proportional Reporting Ratio (PRR) or Bayesian Confidence Propagation Neural Network (BCPNN) flag combinations that occur disproportionately more often than expected (sakaradigital.com). Choose the monitoring schedule separately from the statistical or ML method. EMA recommends that each organisation determine the appropriate EudraVigilance monitoring frequency for each active substance or medicinal product (GVP Module IX). Their limitations include:

- **Structured Data Dependency:** If events are buried in free-text (e.g. narrative descriptions, forum posts, literature), traditional disproportionality misses them unless a human has coded them into structure (sakaradigital.com).
- **Single-Source Focus:** Conventional algorithms typically analyze one database at a time; signals that only emerge by correlating multiple sources (e.g. sales data plus EHR plus spontaneous reports) can be missed.
- **Bias towards Known Patterns:** Disproportionality assumes known event categories. It may under-detect novel or complex interactions (e.g. idiosyncratic ADRs triggered by genetic factors) (sakaradigital.com).

AI-powered detection addresses these gaps through several capabilities (sakaradigital.com):

- **Continuous Surveillance:** Evaluate continuous processing as a potential workflow capability. Test any detection-time benefit prospectively against a comparator using the same data and monitoring schedule; retrospective discrimination metrics do not establish a days-or-weeks advantage. Warner et al. conclude that the true time-saving value across therapeutic areas and drug classes requires further study (Warner et al.).
- **Multi-Source Integration:** Advanced algorithms can simultaneously ingest heterogeneous data. An AI model might fuse ICSRs with electronic health record signals, literature-derived insights, and patient-generated data (social media, registries) (sakaradigital.com). For example, it can correlate a spike in EHR-recorded chest pain with related posts on patient forums, triangulating a signal that neither source alone would reveal.
- **Unstructured Data Mining:** Natural Language Processing (NLP) extracts information from free text. Crucial details in an ICSR narrative (e.g. “patient developed severe neuropathy 4 weeks after dosing”) can be parsed and structured for analysis (sakaradigital.com). Similarly, NLP can monitor unstructured sources: one pilot successfully screened social media for COVID-19 vaccine side effects (link.springer.com).
- **Capturing Complex Relationships:** Deep learning models can learn non-linear patterns and interactions beyond raw counts. As SakaraDigital notes, they can detect subtle shifts (temporal clustering, co-medication interactions) that standard ratios miss (sakaradigital.com). In practice, ensemble ML methods (random forest, gradient boosting) have shown markedly higher accuracy than PRR/ROR on benchmark tasks (link.springer.com).

Evidence of AI Effectiveness

Academic reviews and pilot studies collectively indicate that AI can indeed enhance signal detection, though careful design is needed. A recent open-access review found that “machine learning algorithms generally outperformed traditional frequentist or Bayesian measures of disproportionality” in retrospective analyses (link.springer.com). In several studies, supervised ML classifiers (gradient boosting, random forests) achieved very high ROC-AUC scores (up to ~0.97) for identifying known safety signals while controlling false alarms (link.springer.com). Bae et al. reported GBM AUCs of 0.97 for nivolumab and 0.93 for docetaxel (Bae et al.). In a separate study, Dauner et al. found that all models trained on their balanced training set performed similarly (Dauner et al.).

Dong et al.'s COVID-19 vaccine proof of concept evaluated GPT-3.5 for terminology mapping: it correctly assigned MedDRA preferred terms with 78% accuracy (Dong et al.).

Practical Interpretation: FDA notes that adverse-event reporting data can contain duplicate and incomplete reports (FDA dashboard).

In addition, innovative approaches are emerging. Federated learning – where multiple organizations train a shared model on their private data – is being explored for signal detection. By allowing companies and health systems to collaboratively build AI models without exchanging raw data, federated methods could dramatically improve statistical power without violating privacy (sakaradigital.com). If realized, federated signal surveillance could draw on global datasets in a way unachievable by any single entity, surfacing rare ADRs in smaller markets or patient subgroups.

Overall, **current evidence suggests AI can be a powerful supplement to traditional PV analytics** (link.springer.com) (sakaradigital.com). Machine learning excels at high-volume, repetitive tasks (triage, coding) and at pattern recognition in big data (sakaradigital.com). It is already being applied in practice for these purposes. However, signal detection is safety-critical. The same studies also repeatedly emphasize that performance must be **interpretable and validated in context**. AI's ability to spot a signal should not substitute for human medical judgment; rather, it should augment analysts by presenting prioritized hypotheses for further evaluation (sakaradigital.com).

Implementing AI Governance in Signal Detection

The previous section highlighted AI's capabilities in PV. Equally important is **how** to deploy AI so that its benefits are realized safely. This calls for a comprehensive *AI governance framework* – policies, processes, and oversight to manage AI systems in PV. Both CIOMS WG XIV and industry experts agree that governance must address multiple dimensions: risk management, human oversight, validation, transparency, data ethics, and accountability. Here we analyze these aspects and present structured approaches drawn from recent literature and guidelines.

Core Governance Principles

CIOMS WG XIV distilled several core governance principles – essentially guardrails – for AI in PV (pvpharm.com) (sakaradigital.com). We summarize and align these with broader AI best practices:

- **Patient Safety Primacy (Risk-Proportionate Oversight):** The overarching rule is that AI must *improve or at least not diminish patient safety*. Oversight and rigor should scale with risk: a simple triage bot on non-serious ICSRs can be validated with lighter controls than an AI that directly generates safety signals or label changes. SakaraDigital highlights this "risk-proportionate" approach, noting that controls for automating non-serious case intake differ from those for autonomous signal evaluation (sakaradigital.com). CIOMS similarly emphasizes risk-based oversight and a phasing-in of AI according to criticality (pvpharm.com) (sakaradigital.com).

- **Human-in-the-Loop / Human Oversight:** Qualified PV professionals must remain responsible for decisions. The complexity and sensitivity of the task should determine risk-based human oversight (CIOMS WG XIV authors). Organizations should document mandatory review, escalation, and intervention points for the intended use. WG XIV defines “Human-in-the-Loop” (HITL; humans train or review each output) and “Human-on-the-Loop” (HOTL; humans periodically audit that AI is working correctly) and calls for implementing these as appropriate (pvpharm.com). For example, a pilot might begin with all AI-flagged signals confirmed by safety scientists before acting.
- **Validity and Robustness:** AI models must be fit for purpose, performing reliably on real-world data. This means rigorous validation before deployment and continuous monitoring for drift (sakaradigital.com). Validation should cover range of use cases and inputs, ensuring models properly handle data variability and edge cases (sakaradigital.com). CIOMS expects organizations to demonstrate that their AI consistently meets predefined performance metrics even as data evolves.
- **Transparency and Explainability:** Users must be able to understand or at least audit AI outputs. “Black box” models have limited acceptability in a regulatory context. Whenever possible, organizations should use interpretable models or provide explainability tools that clarify AI reasoning (pvpharm.com) (sakaradigital.com). Even when complex neural networks are used, their decision logic (and any uncertainty) should be documented so that safety experts can challenge or contextualize results (sakaradigital.com). CIOMS and others stress disclosing both *when* AI is used and *why* a given output (e.g. "this case is high-risk") was generated, to build trust and facilitate inspection.
- **Data Privacy and Security:** AI in PV often uses sensitive personal health data (pvpharm.com). Governance must ensure compliance with privacy laws. Protect patient information, including protected health information (PHI), from unauthorized disclosure through an LLM’s training or outputs. Select safeguards, including de-identification where appropriate, for the use case and applicable law. Under HIPAA, covered entities may use or disclose PHI for research with individual authorization or under specified limited circumstances (HHS). CIOMS explicitly warns that generative AI could re-identify patients from “anonymized” PV datasets if misused (pvpharm.com).
- **Fairness and Bias Mitigation:** ML models learn from historical data, which can embed biases (e.g. underreporting in certain countries or demographic groups). Governance must include bias audits and corrective measures (pvpharm.com) (sakaradigital.com). For instance, if social media data skews younger, an AI signal system might miss ADRs prevalent in older patients. Companies should test models on diverse subpopulations and ensure equitable signal sensitivity. The CIOMS report underscores that AI should **not amplify inequities**; fairness checks must be documented.
- **Accountability and Governance Structure:** Ultimately, humans (and organizations) are responsible for AI, not the software itself (pvpharm.com) (sakaradigital.com). Clear roles and accountability frameworks must be defined: Who approves an AI tool? Who signs off on its outputs? Who monitors its compliance? CIOMS makes “governance and accountability” a core principle (not a soft suggestion) (pvpharm.com). Practically, this means high-level oversight committees, model risk teams, and formal procedures for change management.

These principles, synthesized from CIOMS, SakaraDigital and other sources, form the backbone of an AI governance policy for PV. Next, we detail **operationalizing** these concepts in the lifecycle of signal detection AI.

Organizational Governance and Oversight

Translating principles into practice requires specific structures. Expert recommendations typically propose layered governance, as illustrated below:

- **AI PV Steering Committee:** A cross-disciplinary body (PV, data science, IT, quality, regulatory) that provides strategic direction. It approves major deployments, allocates resources, and resolves conflicts. The SakaraDigital blog suggests constituting such a committee to oversee AI strategy (sakaradigital.com).
- **Data Governance Board:** Responsible for the datasets driving AI. It sets standards for data quality, lineage, and bias assessments. For example, this board would review which ICSR fields are used, how missing data is handled, and how training sets are curated (sakaradigital.com).
- **Model Risk Management (MRM) Function:** An independent review team (akin to financial institutions' MRM) that challenges each model's development and controls. MRM ensures no single PV or IT group has unchecked authority. It replicates the risk assessment and validation that the developer did, providing a second opinion (sakaradigital.com).
- **Change Advisory Board:** For computer systems, any significant change (e.g. to AI model version or config) typically undergoes CAB review. The governance model extends this concept: any retraining, parameter tweak or data pipeline change passes through a formal review for regulatory impact (sakaradigital.com).
- **Quality Assurance (QA)/Validation Group:** Works closely with MRM to audit the validation process. In a regulated context, QA ensures the CSV documentation meets standards. For AI, QA must adapt validation docs (see Table **Adapting Validation** later).

By mapping out these roles and committees, an organization establishes *who* makes AI-related decisions and *how*. Documentation of this governance – charters, meeting minutes, decision logs – is itself vital evidence for regulators that AI use is controlled.

AI Model Lifecycle and Validation

A characteristic theme is that AI model deployment is not "set-and-forget." Models degrade as data shifts or as new drug patterns emerge. Governance requires a lifecycle approach (sakaradigital.com):

1. **Foundation Phase (Months 1–6):** Before any model coding, ensure prerequisites are in place. Key activities include:

- **Data Assessment:** Evaluate the quality and completeness of PV data sources. Clean, normalize, and, if needed, link disparate systems (e.g. unify legacy ICSR, EHR and literature data) (sakaradigital.com).
- **Use-case Selection:** Perform risk-benefit analysis to pick candidate use cases. As SakaraDigital advises, prioritize high-volume, routine tasks where error can be tightly controlled, such as coding non-serious ICSRs (sakaradigital.com).
- **Governance Setup:** Formulate the committees and procedures above. Draft AI governance policies (validation plan template, roles/responsibilities matrix).

- **Infrastructure:** Deploy the computing stack (data lakes, model development tools) and define technical standards (e.g. MedDRA/E2B for inputs (sakaradigital.com)).
- **Risk Assessment:** For each planned AI use, identify failure modes (misclassification, bias). Document risk mitigation (e.g. thresholds to trigger human review).

Phase 1 deliverable example: A validated dataset of past ICSRs with coded outcomes that the pilot AI will learn from.

2. **Pilot Phase (Months 4–12):** Begin controlled AI deployment in non-critical settings (sakaradigital.com). SakaraDigital outlines a typical pilot plan:

- **Scope Definition:** E.g. “Automate triage of non-severe, non-expedited ICSRs.” This is an example of a low-risk use-case recommended by AI governance guides (sakaradigital.com).
- **Parallel Operation:** Run the AI model in shadow mode: it scores cases, but humans still make final decisions. This allows measurable comparison (the “ground truth” from human processing vs AI’s suggestions) (sakaradigital.com).
- **Validation and Testing:** Collect performance data (accuracy, recall, time saved). Adjust model hyperparameters and retrain if necessary, under the oversight of QA and MRM.
- **Human Oversight:** Ensure that at defined points in the workflow, humans must verify AI outputs. For instance, a working procedure might say “If the AI flags a serious event, a senior PV physician must review the case.”
- **Documentation:** Record all changes to model parameters, validation results, and any issues. This audit trail is critical for inspection readiness (sakaradigital.com).

Phase 2 deliverable example: Proof-of-concept report showing AI maintained ≥95% concordance with human triage on a test set, along with recommendations for going live.

3. **Expansion Phase (Months 10–24+):** Once pilots demonstrate acceptable performance, scale up.

- **Broaden Use Cases:** Introduce AI to more complex processes (e.g. flagging ICSRs for expedited reporting, supporting signal generation).
- **Integration:** Embed AI systems into existing workflows and databases. For example, an AI module might feed its high-confidence signal alerts directly into the signal management dashboard.
- **Continuous Monitoring:** Implement automated checks for model drift (e.g. input data distribution monitoring) and key performance indicators (are false alarms creeping up over time?) .
- **Retraining Cycles:** Regularly refresh models with new data. Governance protocols should stipulate when retraining is triggered and require re-validation before redeployment.
- **Audit and Inspection Preparedness:** Maintain up-to-date records (model cards, training data snapshots, change logs) in anticipation of regulatory audits.

Table 1 below contrasts traditional computer system validation (CSV) with additional evaluation needed for AI in PV. Model performance evaluation should supplement system requirements and testing while retaining applicable GxP and other standards (FDA–EMA principles).

Validation Aspect	Traditional CSV Approach	AI/ML-Adapted Approach
Requirements Definition	Fixed functional specifications (program does exactly X)	Retain functional and system requirements and their traceability; add predefined model-performance criteria appropriate to the intended use
Testing Methodology	Scripted test cases for specific inputs	Retain functional, integration, security, edge-case, and failure-mode testing; add statistical evaluation on representative hold-out datasets and test the complete workflow, including human interactions
Change Management	Versioned releases with full regression testing	Model versioning with automated drift alerts; re-validation triggered by significant data shifts or architecture changes
Ongoing Compliance	Periodic configuration review	Continuous performance monitoring with automated alerts; routine performance audits (e.g. annual retraining)
Documentation	Detailed test scripts/execution records	Retain requirements, test records, and traceability; add model cards, training-data documentation, bias assessments, performance benchmarks, and audit trails

Table 1. Adapting Computer System Validation for AI in Pharmacovigilance.

In practice, document both the system's required functions and the model's performance criteria. Evaluate hold-out metrics alongside functional, integration, security, and failure-mode tests; include edge cases and human interactions in the complete workflow. Manage changes through the quality system, with risk-based regression testing and model re-evaluation before deployment.

Data Governance and Standards

AI's effectiveness and trustworthiness depend on data. Governing the data lifecycle is as important as the models themselves (sakaradigital.com). Key considerations include:

- **Data Quality and Representativeness:** Ensure training and testing datasets are comprehensive and unbiased. For example, if historical PV data has gaps (some countries under-report, or certain age groups underrepresented), these gaps must be corrected or handled (by weighting or additional data collection) to prevent the model from learning spurious patterns.
- **Data Lineage and Versioning:** Record where input data came from (e.g. which ICSR database versions, literature databases) and how it was preprocessed. This is crucial for traceability if later a signal is questioned. A data governance board should set policies for data version control and archiving.
- **Privacy and De-identification Controls:** HIPAA permits research involving individually identifiable health information under applicable conditions (HHS). For model training, establish the applicable legal basis, permissions, safeguards, and vendor controls in each jurisdiction. Use de-identification where appropriate and feasible, assess residual re-identification risk, and prevent inadvertent disclosure of patient details.
- **Standards Alignment:** Use and enforce recognized coding standards (MedDRA for terms, E2B(R3)/XML for transmissions) to ensure AI systems can interoperate with existing PV tools (sakaradigital.com). The CIOMS/Sakara guidance notes that AI designs should assume interoperability

– e.g. models ingesting AE data must expect MedDRA-coded facets to align across sources.
- **Bias Audits:** Regularly test models for performance across subgroups (age, gender, ethnicity, geography). If the model underperforms systematically for a subgroup, address it via data augmentation or algorithm adjustments (pvpharm.com) (sakaradigital.com).

By instituting these data governance measures, organizations minimize the risk that a model will learn "garbage in." Clean, ethical data are the foundation for any trustworthy AI in PV.

Organizational Action Items

To put principles into practice, experts recommend concrete steps. For example, SakaraDigital lists **eight critical action items** for life sciences companies deploying PV AI (sakaradigital.com):

1. **Map Current PV Processes for AI Opportunities:** Create a detailed inventory of PV workflows (from case intake to signal review) and identify where AI can fit. Focus on high-volume, routine steps (e.g. triage, duplicate checking) with clear human benchmarks (sakaradigital.com).
2. **Establish Accountability Frameworks:** Clearly assign responsibility for each AI component. Document who "owns" the model outputs, who is accountable for patient safety, and how oversight flows. The principle is that while AI assists, a qualified human must be the final responsible person.
3. **Develop AI-Specific Validation Protocols:** As above, build a tailored validation plan that addresses ML peculiarities: data quality checks, performance metrics selection, bias evaluation, etc. This often means writing new SOPs or expanding existing CSV processes to include AI.
4. **Implement Robust Human Oversight:** Design workflows with forced human checks at critical points. For example, an SOP might require a pharmacovigilance expert to review any signal flagged by AI before any regulatory report is submitted.
5. **Document AI Behavior Thoroughly:** Maintain exhaustive documentation for regulators. This includes algorithm design documents, data used, performance results, and decision logs. Notably, "model cards" – concise descriptions of what a model does and its limitations – are increasingly recommended.
6. **Assess and Mitigate Biases:** Incorporate fairness checks into every model's lifecycle. This might involve algorithmic fairness tests, manual review of flagged cases for subgroup bias, or involving patient groups in model evaluation.
7. **Engage Regulatory Authorities Proactively:** When planning to use AI in PV, inform regulators (EMA, FDA, PMDA, etc.) early. Many agencies appreciate transparency and may offer guidance on expectations. Joint workshops and sandboxes (e.g. FDA's AI in clinical tools initiative) are emerging venues for dialogue.
8. **Plan for Regulatory Evolution:** Build flexibility into the AI governance framework. Regulations like the EU AI Act and potential ICH/GMP updates are on the horizon. Organizations should monitor these developments and design their AI processes to adapt (e.g. by modular validation).

These steps, combined with strong leadership support, ensure that AI deployments in signal detection remain constructive. They also help meet CIOMS-specific calls (e.g. CIOMS expects continuous monitoring and regulatory engagement (sakaradigital.com)).

Case Study: Pilot Implementation of AI Triage

While detailed proprietary case studies are rare in public literature, composite scenarios can illustrate best practices. Consider a large pharmaceutical company's pilot of an ML-based ICSR triage system, following the above governance model:

- **Use-Case:** Automatically prioritize incoming ICSRs by seriousness and expectedness, routing "low-risk" cases for standard handling and flagging potential expedited cases.
- **Data Preparation:** The company consolidates its historical 5-year ICSR database, ensuring fields (e.g. event terms, seriousness criteria) are normalized. Any PHI is removed. SME (Subject Matter Expert) teams label a gold-standard subset of reports as "priority" or "non-priority" to train/validate models.
- **Model Development:** A supervised ML classifier is trained to predict the SME labeling. The algorithm selected is a gradient boosting machine, chosen for interpretability and performance. The model explicitly outputs a risk score.
- **Validation:** The ML model is evaluated on a hold-out set. It achieves 96% sensitivity and 90% specificity at a chosen threshold. These metrics meet pre-defined efficacy requirements (≥95% sensitivity) set by the AI governance committee.
- **Parallel Operations:** For 3 months, new ICSRs are processed both by the old manual triage and the new AI system. The outputs are compared: the model successfully matched the human's triage decisions 95% of the time, and in a few instances it caught a serious case that humans had missed (these are audited to refine model inputs).
- **Human Oversight:** During pilot, even if AI marks a case as low priority, a second reviewer spot-checks a random sample. Any mismatches trigger review of the model.
- **Governance Reviews:** At pilot initiation, the AI Steering Committee approved the project with governance framework. Monthly reviews are conducted to examine model drift (sudden changes in input patterns). All model changes are logged in the Model Risk Management function.
- **Outcomes:** The company reports a ~65% reduction in median case processing time for the pilot batch (sakaradigital.com). Based on consistent performance, the Steering Committee approves full-scale rollout of AI triage. Post-pilot, the system is integrated so that human case processors see the AI's priority score next to each ICSR, aiding faster review.

In this (hypothetical) scenario, the company followed recommended steps: careful data governance, performance validation, and defined human checks. Citations support components of this story (e.g. the phased pilot approach (sakaradigital.com) and efficiency gains (sakaradigital.com) (pvpharm.com)). Importantly, any model-flagged signal would undergo further human evaluation as per policy, exemplifying CIOMS's insistence that **humans remain "in control"** (pvpharm.com).

Regulatory Frameworks and Standards

No discussion of AI governance would be complete without recognizing current regulatory stances. While formal regulations on AI in PV are still emerging, agencies have signaled key expectations through publications, pilots, and international collaborations.

- **CIOMS Working Group XIV (2025 report):** CIOMS published its report on Artificial Intelligence in Pharmacovigilance in December 2025 (CIOMS). Companies can use the report as a reference when developing AI governance frameworks.
- **EMA/HMA:** The European Medicines Agency and Heads of Medicines Agencies have been active. In 2025 they initiated a multi-stakeholder workshop on AI (www.ema.europa.eu). EMA's CHMP and CVMP adopted the reflection paper in September 2024 (EMA). Its post-authorisation section states that marketing authorization holders should "validate, monitor and document" their AI/ML models as part of the PV system, especially for signal detection and adverse event surveillance (link.springer.com). EMA's approach emphasizes continuous oversight and integration of AI into existing Good Pharmacovigilance Practices (GVP). Under the EU AI Act, high-risk classification depends on the system's intended purpose and the relevant EU product-safety legislation (European Commission). Assess each PV use case before asserting that high-risk obligations or third-party conformity assessment apply.
- **FDA (U.S.):** CDER's Emerging Drug Safety Technology Program administers Emerging Drug Safety Technology Meetings (EDSTMs) to facilitate discussion and mutual learning about industry's use of emerging technologies in pharmacovigilance (FDA).
- **World Health Organization (WHO):** WHO's *Ethics and governance of artificial intelligence for health* guidance identifies ethical challenges and risks and six consensus principles for AI in health (WHO guidance).
- **Other regulators:** Japan's PMDA has an AI Action Plan (2025) focusing on agency use of AI; presumably it will also look at signals. Canada's Health Canada and Australia's TGA have published conceptual frameworks for AI in healthcare but not PV-specific guidelines yet.

Implications for Compliance: CIOMS and TransCelerate frameworks can help structure governance work, but organizations should assess their documentation and controls against the legal and regulatory requirements applicable to each use case. For experimental technology, EMA offers early interaction through its Innovation Task Force (ITF) (EMA reflection paper).

Discussion and Future Directions

Implementing AI governance in PV signal detection is not a one-time project but an evolving journey. Here we discuss broader implications, challenges, and future trends.

- **Shifting PV Paradigm:** As PVpharm and CIOMS envisage, AI could transition PV from a retrospective reporting discipline to a continuous, predictive science (pvpharm.com). For example, with robust AI, a safety team might monitor real-world signals in real-time (e.g. using hospital EHR feeds) rather than wait for case reports. Such a paradigm shift requires cultural as well as technical change: PV professionals will need AI literacy, focusing more on interpreting AI findings than on manual case processing (pvpharm.com).
- **Workforce and Training:** CIOMS emphasized that PV experts must adapt. Training programs must now include data science fundamentals and ethics. Companies may hire data scientists or upskill existing staff. TransCelerate's approach of combining bio-pharma domain experts with technical teams (e.g. as was done for IGR PV) will likely repeat in AI governance bodies.

- **Bias and Equity:** Despite best efforts, AI may reveal biases. For example, if a model was trained mostly on adult clinical trial data, it may under-detect pediatric signals. Governance must include ongoing bias audits. Transparency principles mean that companies should report known model limitations (e.g. “This signal detector is less reliable on drugs with fewer than 100 ICSRs”).
- **Patient Involvement:** CIOMS’s patient involvement report emphasizes involving patients throughout a medicine’s life, including ongoing monitoring and safe use (CIOMS). Future governance could involve patient representatives in reviewing AI tools (analogous to patient reps on safety committees). This aligns with a general shift toward “patients as partners” in PV, ensuring that patient rights and perspectives shape AI use.
- **Inter-company Collaboration:** Federated learning (as mentioned) could be a game-changer for signal detection. However, it requires industry collaboration on standards. Consortia (like TransCelerate or new alliances) may sponsor federated PV platforms. Regulatory bodies might also facilitate anonymized data sharing frameworks.
- **Advanced Technologies:** Beyond current ML/NLP, generative AI (large language models) will loom large. These can summarize vast literature or generate hypotheses from data, but also hallucinate. Governance must evolve for generative AI specifically. For example, openAI’s GPT might be used to draft PV narratives; oversight would require a new set of tests (e.g. verifying no hallucinated drugs).
- **Regulatory Harmonization:** Expect initiatives like ICH to eventually address AI. Already, ICH’s eWG might expand scope beyond data harmonization to include AI tools. Harmonized global standards will ease trans-national pharmacovigilance. Until then, organizations will try to satisfy the strictest regime (likely EU’s) as a baseline.
- **Ethical and Legal Considerations:** Beyond technical governance, there are bigger questions. If an AI misses a safety signal and patient harm results, liability falls on the MAH. Thus, legal frameworks may evolve to attribute risk. Ethically, the industry and regulators must balance innovation with the precautionary principle; governance frameworks embody this balance.
- **Continuous Improvement:** Feedback loops will refine governance. For instance, after an AI use is audited (say, after an FDA inspection), lessons learned (gaps in documentation, unexpected biases) should revise the corporate policy. Formalizing these lessons (in SOPs or updated CIOMS guides) will make AI governance a living system.

Conclusion

The integration of AI into pharmacovigilance signal detection holds immense promise: accelerating identification of adverse reactions, expanding data horizons, and ultimately protecting patients more effectively (sakaradigital.com) (pvpharm.com). However, this opportunity comes with responsibility. AI models, if unchecked, can introduce new risks that undermine trust in PV. Achieving the promise of AI **requires explicit governance** – a combination of principled guidelines and pragmatic execution.

This report has surveyed the landscape as of 2026: the CIOMS WG XIV report and industry programs like TransCelerate’s autonomous PV initiatives collectively provide a roadmap. Key themes emerged: **Patient safety must remain paramount**, with AI serving as a tool under human supervision; **transparency and accountability** guard against “black box” failures; and **continuous validation** ensures models do not drift or

amplify biases. Organizations must thus institutionalize AI governance: assigning clear roles (Steering/QA committees, Model Risk owners), embedding validation into existing GxP quality systems, and nurturing a culture that scrutinizes rather than blindly trusts AI.

The initial focus will be on “low-hanging fruit” processes (e.g. ICSR triage, coding, literature screening) (sakaradigital.com). But in parallel, innovators envision ambitious futures: predictive safety analytics, real-time population monitoring, even “smart” medical devices that report signals. For each advance, the governance net must adjust: for example, a wearable device generating safety alerts would require IoT security protocols and new privacy rules. The key, as emphasized by CIOMS, is that **holistic governance evolves with technology**. Regulatory authorities and industry must continue collaborating so that guidelines keep pace with innovation.

In conclusion, implementing AI in PV is not an all-or-nothing gamble but a careful ascent. By following the principles and processes outlined by CIOMS and TransCelerate (and supported by academic research), pharmaceutical companies can steadily incorporate AI into their signal detection pipelines. This journey involves not only technical development but also organizational change and ethical vigilance. When done correctly, AI governance transforms signal detection from a reactive, siloed process into a proactive, integrated system – ultimately enhancing the safety net for all medicines and the patients who rely on them.

Key Takeaways: Pharmacovigilance signal detection is poised for an AI-led revolution. Comparative studies reviewed by Warner et al. do not establish that ML algorithms are superior in detection timeliness (Warner et al.). However, to reap these benefits, sponsors must enact robust governance: risk-based oversight, human-involvement, rigorous validation, and ethical data management are non-negotiables (pvpharm.com) (sakaradigital.com). The confluence of TransCelerate’s industry guidance and CIOMS WG XIV’s global framework offers a comprehensive blueprint. By embracing these guidelines and building the requisite organizational structures, the PV community can deploy powerful AI tools **without compromising** patient safety – truly advancing drug safety in the digital age (pvpharm.com) (sakaradigital.com).

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://cioms.ch/working_groups/working-group-xiv-artificial-intelligence-in-pharmacovigilance
2. https://pvpharm.com/navigating-the-ai-revolution-in-pharmacovigilance-insights-from-the-cioms-wg-xiv-report
3. https://sakaradigital.com/blog/ai-powered-pharmacovigilance-signal-detection-3
4. https://link.springer.com/article/10.1007/s40264-022-01164-5
5. https://www.transceleratebiopharmainc.com/initiatives/intelligent-automation-opportunities-in-pharmacovigilance
6. https://www.fda.gov/drugs/science-and-research-drugs/cder-emerging-drug-safety-technology-program-edstp
7. https://dspace.ut.ee/server/api/core/bitstreams/a3721f9e-eb9f-4ecc-9d68-3dc9e33cbbcc/content
8. https://link.springer.com/article/10.1007/s40290-025-00561-2
9. https://www.ema.europa.eu/en/news/ema-fda-set-common-principles-ai-medicine-development-0
10. https://journals.sagepub.com/doi/full/10.1177/20420986231219472
11. https://www.fda.gov/media/180347/download

12. https://cioms.ch/in_media/artificial-intelligence-in-pharmacovigilance-guiding-principles-from-the-cioms-working-group-xiv
13. https://www.transceleratebiopharmainc.com/transforming-pharmacovigilance-solutions-to-solve-regulatory-challenges
14. https://www.ema.europa.eu/en/documents/scientific-guideline/guideline-good-pharmacovigilance-practices-gvp-module-ix-signal-management-rev-1_en.pdf
15. https://www.frontiersin.org/journals/pharmacology/articles/10.3389/fphar.2020.602365/full
16. https://link.springer.com/article/10.1007/s40264-023-01381-6
17. https://www.fda.gov/drugs/fda-adverse-event-monitoring-system-aems/fda-adverse-event-monitoring-system-aems-public-dashboard
18. https://link.springer.com/article/10.1007/s40264-026-01708-z
19. https://www.hhs.gov/hipaa/for-professionals/special-topics/research/index.html
20. https://www.fda.gov/media/189581/download
21. https://www.ema.europa.eu/en/events/hma-ema-multi-stakeholder-workshop-artificial-intelligence-ai
22. https://www.ema.europa.eu/en/about-us/how-we-work/data-regulation-big-data-other-sources/artificial-intelligence
23. https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-use-artificial-intelligence-ai-medicinal-product-lifecycle_en.pdf
24. https://digital-strategy.ec.europa.eu/en/faqs/navigating-ai-act
25. https://www.fda.gov/drugs/cder-emerging-drug-safety-technology-program-edstp/emerging-drug-safety-technology-meetings-edstm
26. https://www.who.int/publications-detail-redirect/9789240037403
27. https://cioms.ch/publications/product/patient-involvement

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore AI enablement, pharma and biotech software development, data engineering and BI, and Veeva services. Contact IntuitionLabs to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: https://intuitionlabs.ai

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.