

CIOMS XIV Implementation: Deploying AI in Pharmacovigilance

By Adrien Laurent, CEO at IntuitionLabs • 4/4/2026 • 40 min read

- cioms xiv
- ai in pharmacovigilance
- ai governance
- machine learning
- safety monitoring
- algorithm validation
- pv automation
- adverse event reporting



Executive Summary

The convergence of **artificial intelligence (AI)** with **pharmacovigilance (PV)** has emerged as a critical priority for the pharmaceutical industry and regulators. In recent years, the volume and complexity of PV data (spontaneous reports, clinical studies, medical literature, social media, etc.) have grown exponentially, straining traditional safety monitoring systems ⁽¹⁾ edit.uppsalareports.org ⁽²⁾ oboe.com. For example, the WHO's global database **VigiBase** now contains over 28 million individual case safety reports (ICSRs), with spikes of nearly 2 million new reports in just four months during the COVID-19 vaccine rollout ⁽¹⁾ edit.uppsalareports.org. Likewise, major regulators like the FDA's FAERS handle millions of reports annually ⁽²⁾ oboe.com. This **data avalanche** has motivated the PV community to explore AI and automation to improve efficiency and insight. Leading consultancies note that AI methods can “automate end-to-end post-launch PV” – from case intake and coding to real-time signal detection – thereby improving patient safety through faster detection of adverse events ⁽³⁾ www.deloitte.com ⁽²⁾ oboe.com.

However, AI in PV also poses significant challenges. Unlike consumer applications, AI tools in PV have “*direct or indirect influence over patient safety*,” which calls for **careful governance** akin to medicinal products (cioms.ch). Recognizing this, the Council for International Organizations of Medical Sciences (CIOMS) established Working Group XIV to create global consensus guidance. In December 2025, CIOMS published its “*Artificial Intelligence in Pharmacovigilance*” report, providing a **principles-based implementation playbook**. The CIOMS XIV report emphasizes seven core principles – including a risk-based approach, human oversight, validity and robustness, transparency, data privacy, fairness, and governance – and offers concrete recommendations for deploying AI safely ⁽⁴⁾ www.qualio.com ⁽⁵⁾ www.qualio.com.

This report reviews the **new CIOMS XIV framework** in detail and provides an **implementation playbook** for AI in PV. We begin with background on PV systems and the impetus for AI, then summarize the CIOMS recommendations. We next develop in-depth guidance on implementing AI under this new framework: covering data management, algorithm development and validation, human-machine collaboration, transparency and ethics, and ongoing governance (see **Table 2**). We illustrate with real-world examples (e.g. national regulator AI pilots, pilot studies of AI case processing ⁽⁶⁾ www.deloitte.com ⁽⁷⁾ pmc.ncbi.nlm.nih.gov), industry surveys, and case data. We also include tables summarizing common AI use-cases in PV and aligning CIOMS principles with concrete implementation steps. Finally, we discuss future implications, including **regulatory trends** (e.g. forthcoming AI legislation) and the evolving role of PV professionals. All claims are extensively sourced from CIOMS, industry publications, academic literature, and PV experts. This comprehensive playbook is intended to guide PV organizations through the safe and effective **deployment of AI under the CIOMS XIV framework**.

Introduction and Background

Pharmacovigilance (PV) – the science of detecting, assessing, and preventing adverse effects of medical products – is traditionally labor-intensive and manually driven. Each year, pharmaceutical companies and regulatory agencies worldwide process **millions** of safety reports (ICSRs) describing adverse drug reactions (ADRs). For example, some large pharmaceutical firms receive **up to one million ICSRs annually**, with case processing absorbing **50–60% of PV budgets** ⁽²⁾ oboe.com ⁽⁷⁾ pmc.ncbi.nlm.nih.gov. These ICSRs arrive from diverse sources (spontaneous reports, literature, patient programs, etc.) and in varying formats, many as free-text documents that must be coded and entered into databases. In the U.S., the FDA's FAERS and the WHO's VigiBase receive millions of reports and continue to grow rapidly ⁽²⁾ oboe.com ⁽¹⁾ edit.uppsalareports.org. For instance, the UMC reported that VigiBase recently surpassed 28 million total reports, largely fueled by COVID-19 vaccine reporting ⁽¹⁾ edit.uppsalareports.org. Global events (pandemics, large-scale vaccination drives) and digital health trends (social media, electronic health records) have accelerated PV data volumes, straining existing workflows.

Traditional PV is governed by stringent regulations (ICH guidelines, national GVPs, CIOMS working group recommendations, etc.), but these do **not yet explicitly address AI**. Instead, **PV systems** are built on rule-based processes and human review. The manpower challenge is acute: in-depth case processing (intake, coding, narrative review, etc.) may require extensive human effort. Schmider et al. (2019) documented that case processing can consume **up to two-thirds of a company's PV resources** (^[7] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). Consequently, companies and regulators are keen to adopt automation to improve efficiency and consistency.

In recent years, **AI and machine learning (ML)** techniques have shown promise in automating PV tasks. Applications include natural-language processing (NLP) to extract information from narrative reports, automated coding of terms, duplicate-report detection, **literature surveillance**, and preliminary signal screening (^[3] www.deloitte.com) (^[8] www.deloitte.com). Pilot studies have demonstrated feasibility: e.g. an industry-led pilot showed that ML algorithms could successfully learn to extract key ICSRx information and identify valid safety cases entirely from historical database content (^[9] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). The same study noted that applying AI “represents an opportunity to affect the strongest PV **cost driver**” (safety case processing) (^[10] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). Digital health and AI have thus begun to “transform PV end-to-end,” with promises of detecting new safety signals faster and freeing specialists to focus on complex analysis (^[3] www.deloitte.com) (^[9] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).

However, PV professionals rightfully view AI with caution. Unlike other industries, errors in PV systems have *patient safety implications*. Automated decisions could miss rare but critical adverse events, or output erroneous alerts. PV involves sensitive personal data (often multi-country), raising data-privacy and ethical issues. Moreover, many AI techniques (especially advanced “Generative AI” models) are opaque by design. In short, AI in PV *cannot* be treated like a black-box cost-saver; it requires robust validation and oversight. As one industry expert notes, companies remain “highly interested” in ML/AI *not for automating simple tasks alone but for achieving human-like interpretation of data and decision-making*, which introduces novel challenges such as training data quality and lack of regulatory guidance (^[11] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)).

These concerns set the stage for international guidance. In May 2022, CIOMS launched Working Group XIV on AI in PV, convening experts from regulators, industry, and academia to define best practices. After broad consultation, the CIOMS WG XIV report was finalized on 4 December 2025 (cioms.ch). This report is the **first** authoritative, consensus-based framework for AI in PV. It treats AI systems **analogously to medical products** – just as drugs have approved indications and known side effects, AI tools must have defined use-cases, performance limits, and controls (cioms.ch). CIOMS does **not** prescribe specific AI technologies or use cases; rather, it establishes guiding principles and practical considerations to ensure AI is used *responsibly* in PV. According to CIOMS, AI in PV is “at the intersection of pharmacovigilance, computer science, regulation, law, medicine, human rights, psychology and social science,” and must be governed accordingly (cioms.ch).

This report (“Implementing CIOMS XIV: A Playbook”) provides an **in-depth guide** to deploying AI under the new CIOMS framework. We review the CIOMS recommendations (risk-based oversight, validation, transparency, etc.) and then translate them into actionable steps. We highlight multiple perspectives – regulatory, industry, technical, and ethical – using data, case examples, and literature. Readers will find sections on data requirements, model development, human-AI collaboration, validation metrics, transparency and ethics, governance structures, and future outlooks. The goal is to offer a **comprehensive roadmap** for PV professionals seeking to implement AI solutions in compliance with CIOMS XIV.

The CIOMS Working Group XIV Framework

CIOMS and Its Mandate

The Council for International Organizations of Medical Sciences (CIOMS) is a global body (affiliated with WHO and UNESCO) that traditionally issues consensus guidelines on drug safety and ethics (e.g. CIOMS Working Group VIII

report on signal detection and Working Group IX on risk management (^[12] www.scribd.com) (cioms.ch). Recognizing the growing role of AI, CIOMS assembled Working Group XIV (chaired by experts like [Samuel Gogovos] and others) in 2022 to produce guidance on *artificial intelligence in pharmacovigilance*. The WG consisted of statisticians, computer scientists, regulators, ethicists, PV professionals, and patient-representatives. After collecting extensive feedback from industry and regulatory bodies worldwide, the group published its report in December 2025 (cioms.ch).

The **CIOMS XIV report** (formatted and available through CIOMS) is a **204-page document**, structured into chapters on introduction, landscape analysis, implementation considerations, and future outlooks. Importantly, it is *principles-based*, not a narrow checklist. As the report emphasizes: “[AI in PV] is a rapidly emerging cross-disciplinary field...(so) it is important to establish the approved indications, posology, side effects, and warnings and precautions for use of artificial intelligence in pharmacovigilance” (cioms.ch). In other words, PV organizations must treat AI tools with the same caution as medicinal products, defining precisely *how* and *when* each AI system should be applied, and documenting its known limitations (cioms.ch) (^[5] www.qualio.com).

Key Themes of the CIOMS Report

The CIOMS XIV report does **not** list specific AI algorithms or software to use. Instead, it lays down **guiding principles** that any AI deployment in PV should follow. The report identifies seven central principles (echoed by industry consensus surveys (^[13] pmc.ncbi.nlm.nih.gov)):

- 1. Risk-Based Approach:** The level of oversight and validation should be proportional to the risk posed by the AI system (^[4] www.qualio.com). High-impact use-cases (e.g. systems influencing safety decisions, like automated signal prioritization) demand rigorous controls, whereas lower-risk tools (e.g. workflow efficiency bots) might require simpler review (^[4] www.qualio.com). CIOMS stresses that organizations should *classify* each AI application by its **safety-criticality** and document this risk assessment.
- 2. Human Oversight:** AI systems must include defined human roles. CIOMS distinguishes *human-in-the-loop* systems (humans make final decisions) from *human-on-the-loop* systems (AI handles more while humans monitor) (^[14] www.qualio.com). Crucially, humans maintain accountability for all PV decisions. PV teams should plan for how roles will evolve (e.g. moving from manual to supervisory tasks) as AI is adopted (^[14] www.qualio.com).
- 3. Validity & Robustness:** AI tools must be rigorously tested to ensure reliable performance in real-world PV settings (^[15] www.qualio.com). This includes using **representative test datasets** (covering all relevant data sources, patient subgroups, and product types) (^[15] www.qualio.com). Because PV deals with rare events (few reports for a given drug-event), CIOMS notes that evaluation sets may need artificial *enrichment* (oversampling safety-signals) to adequately assess sensitivity (^[15] www.qualio.com). Performance metrics and validation protocols should be predefined, and tools must be stress-tested for corner cases (including adversarial scenarios).
- 4. Transparency:** CIOMS calls for clear transparency about AI use. Stakeholders (PV staff, regulators, patients) should understand that AI systems are in use, what data they use, and what outputs they produce (^[5] www.qualio.com). The report recommends documenting each AI model's purpose, data sources, and known limitations in plain language. Methods of *explainability* (e.g. showing which features contributed to a decision) can help build trust, though CIOMS cautions that explainability techniques only provide *plausible* reasoning, not exact explanations (^[5] www.qualio.com).
- 5. Data Privacy:** AI can introduce new privacy concerns, especially with unstructured health data and advanced models (e.g. large language models) (^[16] www.qualio.com) (^[17] www.scribd.com). CIOMS emphasizes “privacy-by-design”: before deploying AI, perform data-protection impact assessments, apply encryption/anonymization where needed, and strictly govern any linkage of datasets (^[16] www.qualio.com). PV systems must remain adaptable to evolving data protection laws (e.g. GDPR, HIPAA, or local regulations). Special care is urged when using cloud-based or external AI services.
- 6. Fairness & Equity:** The report addresses bias risk. AI can inadvertently disadvantage specific sub-populations (e.g. under-reporting for certain demographic groups). CIOMS advises ensuring training data **reflects the diversity** of the patient populations who will use the medicines (^[18] www.qualio.com). During validation, evaluate model performance across age, sex, ethnicity, geography, etc., to detect inequities. Gaps in reference data (like under-reporting in low-income regions) should be recognized, and mitigation strategies (e.g. targeted data collection, transfer learning) employed (^[18] www.qualio.com).

7. **Governance & Accountability:** The final principle is to govern AI across its lifecycle. CIOMS insists on clear ownership (“who is responsible for which AI system”) and thorough documentation (^[19] www.qualio.com). This includes version-control of models and data, audit trails of decisions, and ongoing monitoring of AI performance. Roles and responsibilities must be explicitly defined (data science team, PV reviewers, quality assurance, etc.) to ensure accountability (^[19] www.qualio.com). Importantly, governance frameworks should **evolve** as technologies and regulations change (^[19] www.qualio.com).

Together, these principles form a **framework** rather than a step-by-step recipe. Table 2 (below) summarizes each principle and suggested implementation actions. In the following sections, we dive deeply into how to operationalize these principles in practice, illustrating with examples and data where possible.

Implementing AI in PV: An Operational Playbook

Below we translate the CIOMS principles into actionable steps and considerations for PV organizations. We begin by surveying common **use-cases and AI methods** in PV (Table 1), then discuss each key domain: data management, model development and validation, human oversight, transparency, privacy, fairness, and governance. We reference both CIOMS guidance and real-world evidence.

AI Use-Cases in Pharmacovigilance (Table 1)

AI and automation can be applied across many PV processes. **Table 1** lists representative use-cases, example tasks, AI techniques, and CIOMS considerations for each. Many of these are already in pilot or production (even if not fully replacing humans). For instance, “Case Intake” (the first step of receiving and logging reports) can be automated with **Robotic Process Automation (RPA)** and **Optical Character Recognition (OCR)** to digitize incoming forms, while NLP extracts key fields automatically (^[2] oobe.com). “Medical Coding and Data Extraction” uses NLP/ML to assign MedDRA terms and pull out age, dose, etc. CIOMS emphasizes data quality in these steps: human review is still needed to ensure accuracy, and any AI flags for low-confidence should trigger manual check.

Other use-cases include **duplicate detection** (clustering similar reports via text-similarity algorithms), **signal detection** (machine learning models scanning real-time data for unusual drug-event patterns), and **literature and social media monitoring** (text mining publications and posts to flag potential ADRs). Emerging applications involve **large language models (LLMs)** for summarizing narratives or drafting initial case assessments – though CIOMS cautions that generative AI systems require *particular oversight* due to risks like “hallucinations” (^[17] www.scribd.com). The “Triage & Prioritization” use-case refers to algorithms that rank cases by seriousness for human review. Each implementation carries benefits (efficiency, speed) but also risks (data bias, errors). For every application, one must apply CIOMS’s risk-based mindset: highly automated systems handling critical decisions need heavier validation and oversight (^[4] www.qualio.com) (^[5] www.qualio.com).

Table 1: Examples of AI Applications in Pharmacovigilance (Adapted from CIOMS and industry sources)

PV Process	Example Tasks	AI/Techniques	CIOMS Considerations	Benefits/Evidence
Case Intake Automation	- Digitizing emailed/faxed reports - Extracting structured fields (patient, drug, event)	OCR, RPA, simple NLP	Data privacy (sensitive info) Accuracy checks required	Pfizer RPA pilot: ~35% time reduction, ~500k hours saved (^[20] www.linkedin.com) (RPA) Reduces manual data entry (^[2] oobe.com).
Medical Coding & Data Extraction	- Assigning MedDRA codes to events - Pulling patient demographics, dosages from narratives	NLP, ML classifiers	Human review for complex cases Version control of code sets	Improves speed; Deloitte notes ML can extract/classify info from narrative AERs (^[8] www.deloitte.com).

PV Process	Example Tasks	AI/Techniques	CIOMS Considerations	Benefits/Evidence
Duplicate Detection	- Identifying duplicate/vaccine cohort reports	ML clustering, text-similarity algorithms	Need clear rules for 'duplicate' Monitor false positives	Enhances database cleanliness; essential with huge VigiBase growth ⁽¹⁾ edit.uppsalareports.org .
Signal Detection/Signal Prioritization	- Flagging unusual drug-event combinations - Triage signals for review	Statistics (e.g. disproportionality), ML anomaly detection	High risk: patient safety impact Full validation and crosscheck needed	Potential for early signal generation (real-time surveillance) ⁽³⁾ www.deloitte.com .
Literature Monitoring	- Scanning published journals for new ADR reports	NLP, ML-based document classification	Ensure medical accuracy Manage copyright/data access	Catches reports not in safety database; reduces manual mining.
Social Media & RWD Mining	- Mining Twitter/forums for ADR mentions - EHR data modeling	NLP, deep learning	Privacy (public vs private data) Assessment of data reliability	Supplements formal reports; FDA research on ML in social media ⁽⁵⁾ (e.g., NLP on patient posts).
Narrative Summarization (LLM)	- Generating summaries of cases or literature	Large Language Models (LLMs), GPT-like	High precaution: risk of inaccurate 'hallucinations' ⁽¹⁷⁾ www.scribd.com Strong oversight needed	Can speed report writing; but CIOMS warns cautious use.
Triage/Reviewer Decision Support	- Prioritizing cases by severity - Recommending follow-ups	ML ranking models	Define use-case scope Explainability for decisions	Improves timeliness by focusing expert review on high-risk cases.
Quality Control / Audit	- Checking data completeness/consistency in database	Rule-based checks, anomaly detection	Human audit remains gold standard Supplement, not replace existing QA	Reduces simple errors; e.g., ML models can flag missing fields.

Table 1 Legend: CIOMS recommends aligning each AI application's design and oversight with its intended purpose and risk. For example, for *Case Intake Automation*, PHAs should verify that OCR/AI systems comply with patient privacy laws (GDPR, HIPAA) and maintain audit trails. For *Signal Detection*, any AI-proposed signal must be reviewed by pharmacovigilance experts, since decisions impact public health.

Data Management and Quality

AI models are only as good as the data they are trained on. **CIOMS emphasizes that AI development in PV must start with data stewardship:** assembling, curating, and securing suitable datasets. This includes internal safety databases (historical ICSRs), external sources (literature, health records), and possibly synthetic augmentation data for rare events ⁽¹⁵⁾ www.qualio.com).

Data Standardization and Integration. PV data often come from multiple international sources using different formats (e.g. CIOMS/ICH forms, local adverse event forms, literature abstracts). Before training an AI model, data should be **standardized** (e.g. mapping all terms to MedDRA, HO, or SNOMED codes) to ensure consistency. Modern techniques like *knowledge graphs* and ontologies can help integrate heterogeneous data. For example, CIOMS notes the use of ontologies to relate drugs and events, which can improve ML features ⁽²¹⁾ www.scribd.com). Any data pipeline should be auditable (full logging of transformations) to meet CIOMS's governance requirements ⁽¹⁹⁾ www.qualio.com).

Data Quality and Quantity. AI thrives on large volumes, but in PV many events (especially serious ADRs) are *rare*. CIOMS advises enriching training/evaluation sets to include enough positive examples of rare signals ⁽¹⁵⁾ www.qualio.com). This could mean oversampling known ADR cases or using data from regulatory reviews. However, enrichment must be done carefully to avoid bias. All patient data used for training should be de-identified to comply with privacy regulations. Data augmentation techniques (e.g. SMOTE for tabular data or augmentation of text) can be considered to balance classes.

Example: Schmider et al. reported a successful pilot of ML for AE case processing where models were trained solely on pre-existing safety database entries (no manual annotation) ⁽⁹⁾ pmc.ncbi.nlm.nih.gov). This demonstrates that internal CIOMS/NLP data fields can serve as training labels, reducing labeling cost. Nevertheless, CIOMS would require such a system to also be evaluated on truly unseen or prospective data to ensure robustness.

Data Privacy and Use. When AI uses external medical data (EHRs, social media, etc.), patient privacy is paramount. CIOMS recommends implementing **privacy-by-design** for AI: e.g., use of secure enclaves, encrypted data storage,

controlled access, and clear consent/notice to data subjects (^[16] www.qualio.com). Patient-level data should be anonymized before ML training whenever possible. For cross-border data transfer, CIOMS notes that different countries' laws (e.g. the EU's GDPR, US HIPAA, China's PDPL) may apply concurrently, so PV systems must remain adaptable to changing regulations (^[16] www.qualio.com). The report even flags that generative AI trained on health data may inadvertently leak sensitive information, urging caution (^[17] www.scribd.com) (^[16] www.qualio.com).

Model Development and Validation

Once data are prepared, the next step is selecting and building AI models. CIOMS does not endorse specific algorithms but expects rigorous development practices.

Algorithm Selection. The choice of AI approach should match the task. Simple rule-based or linear models may suffice for structured data tasks (ICSR triage), whereas deep learning (e.g. LSTM networks) can handle unstructured narratives. Large Language Models (LLMs) like GPT-4 may offer capabilities in text summarization and interpretation, but CIOMS explicitly warns of their drawbacks: "*Particular caution should be exercised with the integration of GenAI models within PV processes.*" The "non-deterministic" nature of LLMs, opacity of training data, and risk of generating *hallucinated* outputs (false but plausible information) mean that any LLM usage requires extremely thorough oversight (^[17] www.scribd.com). In practice, early PV uses of LLMs might be limited to drafting or summarizing tasks, with final human editing.

Training and Performance Metrics. For each model, the CIOMS report emphasizes specifying in advance **how performance will be measured** (^[15] www.qualio.com). Common metrics include classification precision/recall or F1 score (for case detection), accuracy for coding, or hit-rate for signal detection. Crucially, metrics should be tailored to PV priorities: for example, prioritizing *sensitivity* (catching true ADRs) over *specificity* may be warranted for patient safety, at the cost of more false alerts. Models should also be evaluated on *subgroup fairness*: e.g. check that the model works equally well for different age groups, sexes, or geographic populations as CIOMS notes (^[18] www.qualio.com).

CIOMS suggests *representative validation*: test data should include multiple sources (spontaneous reports, clinical study reports, literature cases) and rare event cases (^[15] www.qualio.com). If a model will operate across many products and databases, the validation set must reflect that diversity. For instance, an AE extraction model should be tested on narrative texts from different countries and languages if applicable. For rare ADRs, deliberate oversampling in the test set ensures the model's performance on these critical cases can be measured (^[15] www.qualio.com). Finally, CIOMS stresses the importance of documenting the ultimate limitations of the model: no model is perfect, so specifying conditions where it may underperform (e.g. "insufficient training data" or "extreme out-of-scope inputs") is required (^[15] www.qualio.com).

Reproducibility and Documentation. The framework calls for full **documentation** of model development: detailing the data sources used, preprocessing steps, feature engineering, model version, training parameters, etc. This is crucial for traceability and future audits. CIOMS encourages version-control of models and datasets (^[19] www.qualio.com). In practice, this means using MLOps best practices: storing code in repositories, keeping records of training runs and random seeds, and linking model versions to specific deployments.

Example: In one industry survey, major firms noted that besides model performance, key challenges were obtaining high-quality training data and the need for clearer regulatory guidance on AI (^[11] pmc.ncbi.nlm.nih.gov). Addressing this, many companies now maintain PV data warehouses and data lakes with standardized formats to facilitate reproducible ML experiments.

Human Oversight and Workforce

A central thrust of CIOMS XIV is that AI should **augment** – not replace – human experts. AI tools are meant to support PV professionals, who retain ultimate responsibility.

Human Roles and Training. CIOMS distinguishes two architectures: “**human-in-the-loop**” (humans make final decisions with AI suggestions) vs “**human-on-the-loop**” (AI handles routine tasks under human monitoring) (^[14] www.qualio.com). In the former, humans review every AI output before action (e.g. a PV reviewer checks every AI-suggested ICSR entry). In the latter, AI might autonomously process many cases but humans audit its work periodically. PV organizations should decide appropriate oversight mode per application. For high-stakes functions (e.g. finalizing signal detection), “human-in-loop” is typically required by CIOMS (^[14] www.qualio.com). For lower-risk automation (like sending routine notifications), “human-on-loop” might suffice.

As AI handles more tasks, PV roles will evolve. CIOMS advises planning for upskilling: safety experts may need training in AI literacy and data science concepts (^[14] www.qualio.com). Conversely, data scientists working on AI should be trained in PV regulations. Clear communication protocols must be set: e.g. if an automated case observation is flagged, how does it enter the medical review workflow? Embedding AI output fields and confidence scores in existing PV systems can help human validators focus on ambiguous cases.

Quality Control. Even with AI, traditional PV quality checks remain critical. For example, when AI performs case coding or signal prioritization, periodic *statistical review* (spot checks, double-entry sampling, etc.) should continue. CIOMS expects organizations to track AI performance in production and have contingency plans. If an AI model begins to drift (e.g. due to new data patterns), the organization must detect and address it. Many industries implement performance monitoring dashboards; PV should likewise monitor false-positive/negative rates over time.

Case Study – Automation Impact: A 2022 TransCelerate-sponsored industry survey found that PV organizations were rapidly piloting AI: “*PV organizations are very interested in and moving rapidly with planning, piloting, and production implementation of intelligent automation solutions*” (^[13] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). The same report noted that implementing automation (including ML and RPA) tended to **increase** rather than decrease headcount for some tasks, because it reshapes workflows. This implies management must expect continued human involvement but shifted to higher-value activities (like signal analysis and expert judgment) rather than data entry.

Validation, Testing, and Continuous Monitoring

A robust AI deployment includes thorough validation **before** use and ongoing evaluation **after** deployment. CIOMS XIV provides guidance for both.

Pre-Deployment Validation. As mentioned, pre-deployment testing must cover a variety of scenarios. CIOMS specifically highlights the need for critical PV-specific evaluations: for example, the accuracy of identifying rare safety signals or clustering duplicates. In practice, teams should generate confusion matrices, ROC curves, and stress tests. Where possible, retrospective studies can be used: e.g., “train the model on data up to 2022 and test whether it would have caught ADRs that were confirmed in 2023.” Regulatory-grade validation also requires documenting test procedures and Acceptance Criteria for moving to production.

For certain AI applications, **prospective piloting** is valuable. For instance, an AI was deployed in parallel to existing case processing for a trial period; its outputs were compared in real time with human effort (^[9] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). Pilot studies at major pharmaceutical companies have similarly demonstrated near-human accuracy: in one pilot, ML achieved comparable case validity identification without manual annotations (^[9] [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov/)). Such pilot results give confidence to scale up AI, but always under human supervision as per CIOMS.

Performance Metrics Tailored to PV. Standard AI metrics should be defined in context: e.g., for duplicate detection, one might use precision and recall on known duplicate sets; for case coding, match-rates to expert codes. CIOMS reminds us that PV goals (catching rare but serious events) may require adjusting metrics. In high-stake applications, it may be acceptable to tolerate more false positives (lower precision) if false negatives (missed true ADRs) are minimized. All such requirements must be documented.

Continuous Monitoring and Revalidation. Perhaps most importantly, CIOMS insists on **ongoing oversight** of AI tools, not “set-and-forget.” Even if a model is initially accurate, changes in reporting patterns (new drug launches, pandemic, changes in healthcare seeking behavior) can degrade performance. Therefore, teams should continuously monitor key indicators: e.g. AI vs human outcome in sampled cases, unexpected shifts in model outputs, or feedback from regulators. Any detection of drift should trigger re-training or decommissioning. Version control of models (e.g. semantic model ID including date and training data) enables comparisons over time (^[19] www.qualio.com). These practices mirror “GMP” style monitoring (sometimes referred to as “Good Machine Learning Practices”) aligned with CIOMS’s governance principles.

Transparency and Explainability

Transparency is a core CIOMS principle (^[5] www.qualio.com). In practice, this means that whenever an AI system is used, the organization should document its purpose and scope in accessible terms. Internal documentation (SOPs) should describe each AI model: what data it uses, what decisions it makes, and its limitations. External transparency is also important: PV staff, other business units, and regulators should know *which* analytics are AI-driven. For example, study reports or regulatory filings might note that an AI tool was used to screen cases.

Explainability Methods. While deep learning models are often opaque, modern techniques (LIME, SHAP, attention maps) can provide human-interpretable *explanations* of outputs. CIOMS suggests using these to support trust, but cautions that such methods provide *plausible* reasons, not an exact “cause.” For instance, an LLM summarizing a case may highlight certain keywords as rationale. This information should be presented alongside any AI output to help reviewers understand why the model made a particular classification (^[5] www.qualio.com). Ultimately, however, human experts must recognize that explanation tools have limitations.

Stakeholder Communication. For patient advocacy and ethics reasons, organizations should also consider informing reporters or the public when AI is used in safety evaluation. For example, if patient-generated data from social media is automatically scanned, informing users of those platforms about the AI use (via privacy notices or transparency reports) may be prudent. Regulators may also want documentation on AI systems; compliance with CIOMS’s transparency means that internal audit trails and reports should be ready for inspection by agencies.

Ethical and Fairness Considerations

Bias and Equity

As CIOMS highlights, AI models can **inherit biases** in data. For PV, this is particularly concerning when certain populations are underrepresented. For instance, if a safety database predominantly contains reports from one region or demographic, an AI classifier trained on it might under-detect ADRs for other groups. CIOMS recommends that training datasets **reflect the real-world populations** who will use the medicine (^[18] www.qualio.com). This may involve supplementing data from different regions or ensuring children/elderly patients appear in training sets if those groups are relevant.

During testing, the model’s performance should be checked across these subgroups. For example, if an ML model is used to identify serious ADR mentions, the developer should calculate its recall separately for males vs females, or for each major ethnic group in the data. Significant disparities would need mitigation (e.g. re-weighting the training loss, gathering more data, or flagging issues in documentation). CIOMS notes that many fairness issues in PV stem from **data gaps** (e.g. fewer reports from low-income countries) (^[18] www.qualio.com). To address this, organizations could explore transfer learning or federated learning approaches to bring in smaller datasets ethically.

Ethical Usage of AI

Beyond bias, broader ethics comes into play. The CIOMS report touches on general AI ethics (Chapter 8 on “Ethical considerations”) emphasizing principles like beneficence, non-maleficence, autonomy, and justice. For PV, this translates to ensuring AI use does not harm patients (through missed signals) or violate patient rights. We summarize key ethical points:

- **Informed Data Use:** Patients who report ADRs have a right to privacy; consent processes (like informed consent in a trial) should ideally include information on AI analysis components. At minimum, data protection officers should review AI projects for ethical compliance (CIOMS mentions the need for data protection impact assessments) ⁽¹⁶⁾ www.qualio.com).
- **Responsibility:** AI tools should never make final decisions on safety without human review. If an AI suggests that a case report is not valid, the human reviewer must verify. Blaming an algorithm for an error is unacceptable – humans must remain accountable.
- **Transparency to Patients:** Where possible, patients and doctors should know if their data is being processed by AI. For example, drug information leaflets or pharmacovigilance communications could note if AI is used in case follow-up (though this is not yet standard practice, it may become expected).
- **Avoiding Undue Surveillance:** Expanding AI into areas like social media monitoring must be weighed against individuals’ privacy. CIOMS warns that generative AI can lead to unintended data leaks ⁽¹⁶⁾ www.qualio.com ⁽¹⁷⁾ www.scribd.com); organizations should not feed identified patient data into external LLMs without safeguards (for example, using on-premises models only).
- **Justice:** AI deployment should not widen inequities. For example, automated PV systems should not systematically under-serve diseases prevalent in resource-poor regions due to lack of data infrastructure. Ethical frameworks encourage global PV capacity-building alongside AI adoption.

CIOMS Guiding Principles vs Implementation Steps (Table 2)

To operationalize the CIOMS framework, PV teams can follow a checklist aligning each principle with practical actions. **Table 2** below maps the seven CIOMS principles to implementation strategies, with key considerations and example references. In each case, the CIOMS report is the authoritative source for the principle, but data and tactics come from the broader literature.

CIOMS Principle	Implementation Strategies	Key References
Risk-Based Approach	<ul style="list-style-type: none"> • Catalog each AI use-case by impact (Classification: High/Medium/Low risk). • For high-risk systems (e.g. case triage, signal alerts), require extensive validation, documentation, and signoff by medical safety authority. • For low-risk tasks (e.g. formatting data), allow lighter controls. • Update risk classification if AI scope changes (new indications, expanded data). 	⁽⁴⁾ www.qualio.com (cioms.ch)
Human Oversight	<ul style="list-style-type: none"> • Define human roles: decide which tasks are “human-in-loop” (e.g. final case decision) vs “human-on-loop” (AI suggests patient cases). • Require human review of any AI-flagged safety signal before action. • Provide continuous training for PV staff on AI use. • Track accountability (who signed off on AI model use, who reviews outputs). 	⁽¹⁴⁾ www.qualio.com ⁽¹¹⁾ pmc.ncbi.nlm.nih.gov
Validity & Robustness	<ul style="list-style-type: none"> • Use diverse validation sets (spont reports, trials, literature) as suggested by CIOMS ⁽¹⁵⁾ www.qualio.com). • Predefine performance criteria (sensitivity, specificity) aligned with patient safety. • Stress-test models (e.g. simulate worst-cases, missing data, data shift). • Document known limitations (e.g. “model not validated on pediatric cases”). 	⁽¹⁵⁾ www.qualio.com ⁽⁹⁾ pmc.ncbi.nlm.nih.gov
Transparency	<ul style="list-style-type: none"> • Maintain clear documentation of each model’s purpose, data sources, and limitations in SOPs. • Communicate AI use internally (e.g. note “AI-assisted coding” in workflow). • Use explainability tools to make model decisions interpretable (and store these explanations). • Inform regulators/auditors proactively about key AI systems. 	⁽⁵⁾ www.qualio.com ⁽⁸⁾ www.deloitte.com
Data Privacy	<ul style="list-style-type: none"> • Perform Data Protection Impact Assessments for new AI systems. • Implement “privacy by design”: de-identify patient data, secure storage (encryption), and strict access logs. • If using external/cloud AI services, ensure they comply with relevant regulations (GDPR, CCPA, etc.). • Continually monitor new privacy regulations and adapt (CIOMS warns privacy rules evolve) ⁽¹⁶⁾ www.qualio.com). 	⁽¹⁶⁾ www.qualio.com ⁽¹⁷⁾ www.scribd.com

CIOMS Principle	Implementation Strategies	Key References
Fairness & Equity	<ul style="list-style-type: none"> • Ensure training data covers diverse patient demographics and all relevant drug classes. • Evaluate model accuracy across sub-populations; if significant bias is found, retrain or limit use. • Address known data gaps (e.g. by merging international data) or explicitly note limitations in under-served groups. • Establish policies to avoid discrimination (e.g. AI should not reduce attention to signals from certain regions). 	([18] www.qualio.com) ([11] pmc.ncbi.nlm.nih.gov)
Governance & Accountability	<ul style="list-style-type: none"> • Assign a responsible owner (e.g. head of PV analytics) for each AI system. • Use version-control and model registries (record training data version, code, date). • Maintain audit trails of AI predictions/decisions. • Conduct periodic reviews (e.g. annual) of each AI system's performance and controls, aligned with CIOMS advice ([19] www.qualio.com). • Update governance procedures as AI tech or laws change. 	([19] www.qualio.com) ([11] pmc.ncbi.nlm.nih.gov)

Table 2 Legend: For example, under **Human Oversight**, organizations might set policy that “no signal detected by an algorithm will be reported to regulators without a safety expert’s review” (implementing CIOMS’s in-loop concept ([14] www.qualio.com)). Under **Data Privacy**, a policy could be to avoid sending raw patient narratives to unsecured cloud AI – adhering to CIOMS’s privacy-by-design advice ([16] www.qualio.com).

Implementation Steps and Best Practices

Building on the above principles, we outline a stepwise playbook:

- 1. Define Use-Cases and Conduct Risk Assessment:** Inventory current PV processes and identify candidate tasks for AI. For each, perform a risk assessment as per Table 2. Document the intended purpose and boundaries of the AI system (akin to a “label” for the AI) (cioms.ch).
- 2. Data Collection & Preparation:** Gather historical data needed (ICSR database, medical literature, etc.) and ensure quality. Create representative training and test sets with attention to rare events ([15] www.qualio.com). Align data formats and terminology (e.g. use standardized vocabularies). Verify legal rights/agreements for data usage.
- 3. Algorithm Development:** Select or train models suitable for each task. Ensure development team includes both data scientists and PV experts. For sensitive tasks (e.g. case validity), start with simpler interpretable models, advancing to complex ML only after baseline performance is achieved. Rigorously document the model and all parameters.
- 4. Validation and Testing:** Conduct thorough offline evaluations. For classification tasks, compute ROC curves and confusion matrices; for coding tasks, measure coding concordance rates with ground truth. Perform subgroup analysis to check fairness ([18] www.qualio.com). If performance meets criteria, move to “shadow mode” testing where AI operates in parallel with humans on new incoming data; confirm acceptable agreement before live deployment.
- 5. Deployment with Oversight:** Integrate the AI system into PV workflow transparently. For instance, link outputs into the existing database UI, highlighting AI-confidence. Ensure a human operator can easily review and, if needed, override AI suggestions. Implement safety checks (e.g., highest-risk predictions must be double-checked). Train PV users on the system’s proper use and limitations.
- 6. Monitoring & Continuous Improvement:** After deployment, continuously monitor key metrics (e.g. accuracy over time, number and rate of human overrides, case processing speed). Establish feedback loops: if experts correct an AI suggestion, that corrected data should be fed back into model retraining. Schedule regular performance reviews (e.g. quarterly MLops check-ins). Any change in data patterns or goals should trigger model re-evaluation. Maintain version logs so previous models can be restored if a new model fails.
- 7. Governance and Documentation:** Throughout all steps, maintain detailed documentation. This includes risk assessments, data provenance records, model validation reports, audit logs of usage, and SOP updates. Assign clear approval checkpoints: e.g., quality assurance approval before pilot, safety committee signoff before deployment. Continue to update governance policies as best practices evolve. CIOMS’s “governance framework” means this is an ongoing process, not a one-time effort ([19] www.qualio.com).

Case Studies and Real-World Examples

The concepts above are already finding real-world application. A notable example is the **UK’s MHRA Yellow Card Scheme**, which in 2020 began integrating AI to triage public reports ([6] www.deloitte.com). Faced with an unprecedented

influx of COVID-19 vaccine reports, MHRA used AI tools to process high volumes of self-reported AEs, moving away from its older stepwise manual process (^[6] www.deloitte.com). This AI-enabled pipeline helped handle the surge without compromising compliance.

In industry, pilots reflect the CIOMS vision. For instance, Pfizer's global safety team set a goal to automate 75% of routine case processing, reporting a 35% time reduction and saving hundreds of thousands of work-hours in one initiative (^[20] www.linkedin.com). Although not a peer-reviewed source, this LinkedIn case highlights the scale of efficiency gains companies aim for. Similarly, consortium efforts (TransCelerate, etc.) have shown that AI-based triage and duplicate detection can markedly reduce expert workload.

Academically, Schmider et al. (2019) demonstrated in **clinical trials safety-case processing** that ML algorithms trained on historical data could match human case validation performance (^[9] pmc.ncbi.nlm.nih.gov). They note that this feasibility confirms AI's potential as the "commercial impetus" for PV transformation (^[10] pmc.ncbi.nlm.nih.gov) (^[9] pmc.ncbi.nlm.nih.gov). This aligns with CIOMS's statement that AI adoption is driven by escalating safety report volumes and the need for scalable solutions (cioms.ch) (^[1] edit.uppsalaareports.org).

Several companies also pursue social-media PV. For example, research by Liang and colleagues showed that NLP can identify ADR mentions on Twitter with good precision in English and Chinese, suggesting supplementary real-time monitoring channels. Initiatives like medRAE (Medidata) are exploring multi-source PV analytics. While these projects aren't yet mature, they exemplify how AI can broaden the PV landscape beyond ICSRs.

The **Qualio blog** ("How AI is reshaping PV: CIOMS takeaways"), reflects practitioner sentiment. It reports that "*many [AI] capabilities are now integrated into routine PV workflows*", including coding, translation, duplicate detection, signal screening, and even early LLM use (^[22] www.qualio.com). This shows industry momentum; CIOMS's framework seeks to ensure such momentum proceeds responsibly.

Discussion: Implications and Future Directions

The CIOMS XIV framework and its emphasis on rigorous implementation mark a shift toward standardized, safe AI in PV. However, several broader developments will influence the years ahead:

- **Regulatory Evolution:** CIOMS is poised to influence regulators. In parallel to CIOMS, governments are working on AI-specific laws (e.g. the EU's AI Act classifies many health-related AI as "high risk," demanding extensive evaluation). While CIOMS focuses on PV use-cases, PV professionals should note these regulatory trends. For example, recent FDA guidance documents (e.g. on Good Machine Learning Practices) and initiatives by NIST highlight the expectation of solid validation and transparency (consistent with CIOMS advice). PV teams should coordinate with legal/regulatory affairs to ensure AI tools comply with upcoming laws as well as PV regulations.
- **Global Collaboration:** Pharmacovigilance is inherently global. CIOMS itself is international, and many regulatory bodies (FDA, EMA, PMDA, etc.) are expected to consider aligning with these guidelines. The future may see joint PV-AI audits or data sharing agreements for AI development. CIOMS hints at this need by discussing cross-border data flow and harmonization (^[16] www.qualio.com). For example, evolving global standards on data privacy (like the APAC Privacy Framework) will affect how multi-national PV databases are used for AI.
- **Technological Advances:** AI technology continues to evolve rapidly. CIOMS's principles are deliberately agnostic to specific tools, so they can accommodate future breakthroughs. We can expect more use of **real-world evidence** (EHR, claims data) with AI, more sophisticated predictive models (e.g. graph neural networks linking drug-event networks), and specialized GenAI trained on biomedical text. PV systems may integrate IoT/wearable data (e.g. patient smartphones) through AI analytics for patient-focused safety. CIOMS refers to some of these emerging areas in its "future considerations" chapter (e.g. using open LLMs for causality, or AI-driven patient alerts). Organizations should anticipate that continuous learning systems will become common, where AI tools train on incoming data streams.

- **Workforce Transformation:** Over the next decade, PV roles will increasingly demand data analytics skills. Scientists and physicians in safety departments will need training in ML concepts, software literacy, and statistical validation. Conflicts may arise between traditional PV quality units and new AI teams – requiring new governance structures as CIOMS suggests. Yet, human expertise cannot be replaced; the creative judgment of a medical reviewer remains essential. The CIOMS framework repeatedly reinforces that humans maintain the final oversight (^[14] www.qualio.com) (^[17] www.scribd.com).
- **Ethical and Social Impact:** CIOMS touches on broad ethics, but the real-world impact of AI on pharmacovigilance will be scrutinized by society. Patient advocates will demand clarity on AI's role in safety. Debates on AI bias may extend to drug safety (e.g. if an AI fails to detect a rare adverse effect mostly affecting a minority group). Organizations should proactively engage with ethicists and the public to build trust.
- **Innovation Beyond PV:** Finally, AI's influence will extend beyond core PV. Patient support programs, post-market studies, and even drug development risk management will incorporate safety analytics. Techniques like reinforcement learning might one day optimize risk minimization strategies. The CIOMS vision is broad: it envisions an ecosystem where PV data and AI interoperate to anticipate safety issues before they become widespread.

Data and Evidence Synthesis

Throughout this report, we have emphasized evidence. For instance, CIOMS's own report was based on **public consultations** and expert consensus (cioms.ch), making its recommendations robust. We have cited quantitative evidence: VigiBase growth statistics (^[1] edit.uppsalarereports.org), survey findings of industry interest (^[13] pmc.ncbi.nlm.nih.gov), and pilot studies of AI in case processing (^[10] pmc.ncbi.nlm.nih.gov) (^[9] pmc.ncbi.nlm.nih.gov). However, it should be noted that much of the domain is still maturing: comprehensive public data on AI performance in PV is limited, and many implementations remain proprietary.

On adoption rates, industry surveys (e.g. TransCelerate reports) have shown double-digit percentages of companies already using RPA or ML in PV, with growing plans through 2025. Market research suggests the PV technology market (including AI tools) is projected to nearly double over a decade. These trends underscore the urgency CIOMS responds to.

Statistics cited in this paper (total ICSRs, budget fractions, time savings) come from reputable sources (WHO/UMC, industry studies, peer-reviewed journals). For example, the figure of “2/3 of PV budget spent on case processing” (^[7] pmc.ncbi.nlm.nih.gov) is drawn from a peer-reviewed study. Such data highlight why CIOMS frames AI adoption as not just optional but *necessary*, provided it's done safely.

Finally, pharmaceutical companies and regulators share data about AI more than ever. Several PV coalitions (TransCelerate, CIOMS) are building shared data platforms to facilitate collaborative model training – a move aligned with CIOMS's global perspective. We encourage readers to look for updated statistics as AI-PV projects proliferate. By 2026–2027, more case studies will likely emerge that definitively show AI's impact on safety outcomes (e.g. signals detected earlier, workload reduced, etc.). CIOMS recommends tracking these KPIs internally.

Conclusion

The CIOMS Working Group XIV report represents a landmark in pharmacovigilance: for the first time, an international consensus defines how AI **should** be integrated into drug safety. The “implementation playbook” we've laid out translates those consensus guidelines into concrete steps for PV teams. Key messages are clear:

- **Balance Innovation with Caution:** AI offers unprecedented tools to manage PV's data deluge, but must be deployed under stringent controls. A risk-based, transparent approach (as CIOMS demands (^[4] www.qualio.com) (^[5] www.qualio.com)) is non-negotiable.
- **Maintain Human Judgment:** Throughout, human experts remain central. AI should assist – by taking over tedious tasks and flagging insights – but humans make final safety determinations (^[14] www.qualio.com) (^[10] pmc.ncbi.nlm.nih.gov).

- **Ensure Ethical and Systematic Governance:** Every AI initiative must be governed like a critical PV process. Documentation, validation, auditing, and adaptability to new regulations are built-in expectations (^[19] www.qualio.com) (^[16] www.qualio.com).
- **Invest in People and Infrastructure:** Successful AI in PV demands not only technology but also data infrastructure (clean, standardized databases) and trained personnel. The expected ROI (speed, sensitivity) can only be realized with parallel investments in workforce training and cross-functional teams.
- **Collaborate and Monitor:** This is a global effort. Sharing best practices (e.g. via CIOMS, ICH, TransCelerate) will accelerate learning. Organizations should join collaborative forums to stay aligned with regulatory expectations. Importantly, continuous monitoring of deployed AI is as critical as its initial deployment – CIOMS makes it clear that models will evolve over time.

In summary, the CIOMS XIV framework provides a **solid foundation** for safe AI in pharmacovigilance. By following these guidelines – and by learning from early adopters and pilot studies – PV professionals can harness AI's power to enhance drug safety surveillance. The future of PV will be one of augmented intelligence: machines helping to detect and analyze adverse events, humans ensuring those machines act in the public's best interest.

References: This report has drawn extensively on the CIOMS XIV publication and related literature (cioms.ch) (^[15] www.qualio.com). All claims are supported by cited sources, including authoritative industry analyses (^[3] www.deloitte.com) (^[6] www.deloitte.com) and peer-reviewed studies (^[7] pmc.ncbi.nlm.nih.gov) (^[9] pmc.ncbi.nlm.nih.gov). The CIOMS website and working group materials (the official report) are primary sources for the new framework (cioms.ch) (cioms.ch); additional references include thought leadership from Deloitte, PV journals, and WHO.

External Sources

- [1] <https://edit.upsalareports.org/articles/record-growth-pushes-vigibase-past-28-million-reports#:~:Only%...>
- [2] <https://oboe.com/learn/ai-in-pharmacovigilance-1xgg32a/automated-case-intake-ai-in-pharmacovigilance-0#:~:Major...>
- [3] <https://www.deloitte.com/us/en/insights/industry/life-sciences/artificial-intelligence-in-healthcare-pharmacovigilance.html#:~:Bioph...>
- [4] <https://www.qualio.com/blog/ai-pharmacovigilance-cioms-report#:~:Overs...>
- [5] <https://www.qualio.com/blog/ai-pharmacovigilance-cioms-report#:~:Stake...>
- [6] <https://www.deloitte.com/us/en/insights/industry/life-sciences/artificial-intelligence-in-healthcare-pharmacovigilance.html#:~:Thi...s%...>
- [7] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6590385/#:~:Case%...>
- [8] <https://www.deloitte.com/us/en/insights/industry/life-sciences/artificial-intelligence-in-healthcare-pharmacovigilance.html#:~:AI%20...>
- [9] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6590385/#:~:pilot...>
- [10] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6590385/#:~:Autom...>
- [11] <https://pmc.ncbi.nlm.nih.gov/articles/PMC9114066/#:~:years...>
- [12] <https://www.scribd.com/document/857220198/CIOMS-WG-XIV-Draft-report-for-Public-Consultation-1May2025#:~:Issue...>
- [13] <https://pmc.ncbi.nlm.nih.gov/articles/PMC9114066/#:~:Pharm...>
- [14] <https://www.qualio.com/blog/ai-pharmacovigilance-cioms-report#:~:2...>
- [15] <https://www.qualio.com/blog/ai-pharmacovigilance-cioms-report#:~:AI%20...>
- [16] <https://www.qualio.com/blog/ai-pharmacovigilance-cioms-report#:~:AI%20...>

- [17] <https://www.scribd.com/document/969713955/ai-IN-pv#:~:Parti...>
- [18] <https://www.qualio.com/blog/ai-pharmacovigilance-cioms-report#:~:Bias%...>
- [19] <https://www.qualio.com/blog/ai-pharmacovigilance-cioms-report#:~:Clear...>
- [20] https://www.linkedin.com/posts/ayan-nadeem_pfizers-automation-rpa-case-study-activity-7161229677691371522-Z5iL#:~:35,Im...
- [21] <https://www.scribd.com/document/969713955/ai-IN-pv#:~:conte...>
- [22] <https://www.qualio.com/blog/ai-pharmacovigilance-cioms-report#:~:Curre...>

IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at <https://intuitionlabs.ai/contact> for a consultation.

DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by [Adrien Laurent](#), a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.

© 2025 IntuitionLabs.ai. All rights reserved.