ISSN:2306-6091

International Journal of Pharmaceuticals and Health care Research (IJPHR)

IJPHR|Vol.14|Issue 1|Jan-Mar-2026

www.ijphr.com

DOI: https://doi.org/10.61096/ijphr.v14.iss1.2026.1-7

Review

Artificial Intelligence–Enabled Pharmacovigilance: A New Paradigm in Drug Safety Surveillance

V.Ravali*[1], B. Nagamaheswari, [2]

[1]*Assistant professor, Department of pharmaceutics, Stpauls College of pharmacy, Hyderabad, Telangana, 501506.*
[2]*Assistant professor, Department of pharmaceutical chemistry, Vignan College of pharmacy, Hyderabad, Telangana.*

*Author of Correspondence: V.Ravali
Email: ravaiveldanda@gmail.com.

Published on: 05.01.2026

Published by:
Futuristic Publications

Abstract

Pharmacovigilance (PV) has historically relied on manual intake and clinical review of spontaneous adverse event reports, constrained by under-reporting, variable data quality, duplicate cases, and operational latency. Over the last decade, the digitalization of Individual Case Safety Reports (ICSRs) and the maturation of machine learning (ML), natural language processing (NLP), and "augmented intelligence" workflows have enabled a shift from document-centric, manual processing toward data-centric, semi-automated surveillance. This review synthesizes the evolution from manual safety reporting to AI-enabled PV across the ICSR lifecycle: intake and triage, data extraction and normalization, coding (MedDRA/WHO Drug), case processing and medical assessment support, duplicate detection, signal detection and validation, and regulatory submission. We summarize the enabling regulatory/technical standards (e.g., ICH E2B (R3), EMA GVP), global surveillance ecosystems (FAERS, EudraVigilance, VigiBase), and contemporary guidance (e.g., CIOMS Working Group XIV) that shape responsible AI adoption. We present publication-ready tables describing (i) the PV technology timeline, (ii) AI methods mapped to PV tasks with validation metrics and risks, and (iii) an implementation governance checklist. A figure proposes an end-to-end AI-enabled PV operating model with human oversight and audit ability. We conclude that the most durable value arises from targeted automation (coding, extraction, prioritization) combined with robust governance data quality management, bias monitoring, explain ability, and continuous performance verification rather than full replacement of expert judgment.

Keywords: pharmacovigilance; ICSR; adverse event; NLP; machine learning; signal detection; E2B(R3); FAERS; EudraVigilance; VigiBase; CIOMS.

1. INTRODUCTION

1.1 From "paper cases" to data ecosystems

Pharmacovigilance aims to detect, assess, understand, and prevent adverse effects or other medicine-related problems. For decades, PV operations were anchored in manual workflows: paper forms and narrative reports reviewed by safety professionals who extracted key fields (patient, suspect drug, reaction, and outcome), coded medical concepts, assessed seriousness and expectedness, and prepared regulatory submissions. This model delivered important public health gains but struggled under expanding product portfolios, global reporting obligations, and exploding data volume from post-marketing use.

The modern PV environment is now defined by large, heterogeneous data streams and formalized reporting standards. At the center is the ICSR—standardized safety case data intended for electronic exchange. ICH E2B (R3) provides the ICSR specification and implementation artifacts that enable consistent transmission of structured safety data between marketing authorization holders (MAHs), regulators, and partners [1].

1.2 Global surveillance systems and why scale changed everything

Global safety surveillance is supported by large repositories and national/regional systems:

- **FAERS (US FDA)** supports post-marketing safety surveillance and provides a public dashboard for querying adverse event reports [2].
- **EudraVigilance (EU/EEA)** is EMA's system for managing and analyzing suspected adverse reactions, supporting early detection of potential safety issues [3].
- **VigiBase (WHO Programmed for International Drug Monitoring)** is a global database of ICSRs maintained by the Uppsala Monitoring Centre (UMC) on behalf of WHO and member states; it contains tens of millions of reports and supports signal identification [4].

In parallel, "active surveillance" capabilities using routinely collected healthcare data (claims/EHR) complement spontaneous reporting. The **FDA Sentinel Initiative** is explicitly positioned as a national electronic system for proactive post-market monitoring, launched following FDAAA 2007 and complementing FDA's broader safety monitoring [5, 6].

1.3 Why AI now: operational pressure plus technical feasibility

Three forces converged to accelerate AI in PV:

1. **Volume and velocity:** More reports, more products, more geographies, and increasing expectations for timeliness.
2. **Unstructured narratives:** Many clinically important details live in free text (medical history, temporality, DE challenge/re challenge, and confounders).
3. **Mature ML/NLP tooling:** Named entity recognition, document classification, concept normalization, and deep learning made reliable extraction, prioritization, and coding increasingly feasible.

Recent PV literature increasingly frames the goal as **augmentation**—automating high-throughput, repetitive steps while keeping clinical accountability and regulatory compliance under human oversight [7–9].

2. Methods (PRISMA-style for a narrative / semi-systematic review)

2.1 Review design

We performed a narrative review with semi-systematic elements aligned to PRISMA principles (transparent search strategy, documented screening logic, and structured synthesis). The review emphasizes operational PV tasks where AI is applied: ICSR intake/triage, data extraction, MedDRA/WHODrug coding, duplicate detection, case processing productivity/quality, signal detection, and integration with active surveillance.

2.2 Information sources

We targeted peer-reviewed literature and authoritative guidance from regulators/standards bodies and PV organizations. Sources included:

- Pub Med/Medline (AI/NLP in pharmacovigilance; ICSR processing; social media PV)
- Regulatory and standards sites: ICH (E2B(R3), MedDRA), EMA (GVP, EudraVigilance), FDA (FAERS, Sentinel)
- PV organizations and guidance: WHO-UMC (VigiBase), CIOMS (AI in PV)

2.3 Example search strategy (adaptable)

A representative Boolean query (Pub Med):

- ("pharmacovigilance" OR "drug safety" OR ICSR OR "adverse event reporting") AND ("artificial intelligence" OR "machine learning" OR "natural language processing" OR "deep learning" OR

"large language model" OR automation OR "augmented intelligence") filters: English; 2014–2025; humans.

2.4 Eligibility criteria

Inclusion:

(i) AI/ML/NLP methods applied to PV tasks; (ii) evaluations of automation in case processing/coding; (iii) signal detection methods using spontaneous reports or complementary real-world data; (iv) authoritative guidance (ICH/EMA/FDA/WHO/CIOMS).

Exclusion:

(ii) Pure preclinical toxicology AI (unless connected to PV), non-medicine safety domains, editorials without methods/operational relevance (except landmark guidance).

2.5 Screening and synthesis

Titles/abstracts were screened for PV task relevance, then full texts were evaluated for:

(i) clear task definition,

(ii) Data source and labeling approach,

(iii) Evaluation metrics,

(iv) Risks/limitations (bias, generalizability), and

(v) Operational integration considerations. Findings were synthesized thematically by ICSR lifecycle stage and surveillance objective.

3. Results and Thematic Synthesis

3.1 Manual PV: strengths and structural limitations

Manual PV processing provides nuanced clinical interpretation—especially for complex cases involving poly pharmacy, co morbidity, and ambiguous temporality. However, manual processing is vulnerable to:

- **Inconsistent extraction and coding** across processors and vendors
- **Duplicate cases** and fragmented follow-ups
- **Latency** from intake to signal-relevant aggregation
- **Quality variation** driven by source heterogeneity (consumers vs. HCP vs. literature)

Regulators explicitly caution that spontaneous report repositories support signal generation but do not prove causality; FAERS documentation emphasizes that the presence of reports does not mean a product caused the event and that FAERS data is not a direct indicator of product safety profile [2].

3.2 Digitization and standardization: the substrate AI needs

AI performance depends on standard data structures and controlled terminologies. Key pillars include:

- **ICSR electronic transmission (ICH E2B (R3))** enabling structured fields and consistent exchange [1].
- **Good Pharmacovigilance Practices (EMA GVP Module VI)** guiding collection, management, and submission of suspected adverse reaction reports in the EU framework [10]
- **MedDRA** as a standardized medical terminology enabling consistent adverse event coding and international information sharing [11].

Without these, AI outputs are hard to validate, compare, and audit.

4. AI across the ICSR Lifecycle: where it works, where it breaks

4.1 Intake, triage, and case prioritization

Objective: rapidly identify valid cases, prioritize serious/expedited reports, and route to appropriate workflows (literature, solicited programs, special situations).

AI approaches: document classification, language detection, seriousness prediction, "valid ICSR" criteria extraction, and priority scoring.

Value: reduced cycle time and improved queue management, especially during surges (e.g., new launches, safety crises).

Risks: false negatives (missing expedited cases), distribution shift across regions/languages, and over-reliance on incomplete narratives.

4.2 Narrative text extraction and normalization (NLP)

NLP is used to extract: suspect drug(s), reaction(s), onset dates, dose, indications, medical history, outcomes, and reporter type from free text. Typical architectures include transformer-based models with entity recognition plus normalization to dictionaries (MedDRA for events; WHO Drug for products).

Common failure modes: abbreviations, misspellings, multilingual narratives, implicit temporality, and context (negation: "no rash"). Robust performance requires curate training sets and continuous monitoring.

4.3 Automated coding (MedDRA/WHO Drug) and productivity gains

Coding is a high-volume bottleneck well-suited for augmentation. WHO Drug Koda, evaluated on Vigi Base adverse event reports, illustrates the operational premise: automate high-confidence coding while deferring uncertain cases to humans, improving consistency and throughput [12].

4.4 Case processing augmentation: deep learning for ICSR workflows

A frequently cited example is deep-learning-based cognitive services applied to ICSR processing, designed to extract key characteristics and support real-world case workflows rather than purely academic benchmarks [13].

What matters operationally: not just F1 score, but end-to-end impact—touch time reduction, rework rates, audit findings, and measurable quality improvements.

4.5 Duplicate detection and case linkage

Duplicate reporting is endemic: the same clinical event may be reported by consumers, HCPs, partners, and literature, sometimes with follow-ups. AI can help with probabilistic matching using patient/event/drug similarity and narrative embedding.

Key governance requirement: explainable match rationales and conservative thresholds to avoid incorrect merges.

4.6 Signal detection: from dis proportionality to hybrid ML

Traditional statistical approaches (e.g., disproportionality analysis) remain central for spontaneous report signal detection, but AI increasingly augments:

- feature generation from narratives,
- stratified signal detection (age/sex/region),
- prioritization and triage of candidate signals,
- Literature and label change intelligence.

Active surveillance systems (e.g., Sentinel) complement spontaneous reporting by enabling targeted analyses in healthcare data networks, supporting hypothesis testing and rapid assessment [5].

4.7 Social media and "digital PV": promise with persistent noise

Mining social media for adverse drug reaction (ADR) mentions has a long research history; classic work demonstrated sequence labeling approaches for ADR extraction from informal text [14]. However, social media remains challenged by confounding, unverifiable product exposure, duplicate narratives, and shifting platform access policies. Reviews emphasize the need to filter unsupported claims and triangulate with trusted sources [15].

5. Governance, Regulation, and "Responsible AI" in PV

5.1 Why PV is a high-stakes AI domain

PV outputs affect labeling, risk minimization, and benefit-risk decisions, making accuracy, auditability, and traceability essential. Regulatory compliance also demands that organizations can explain how case data were processed and submitted.

5.2 CIOMS Working Group XIV: emerging expectations

CIOMS has explicitly focused on principles and guidance for AI/augmented intelligence in PV, framing AI as a cross-disciplinary domain requiring careful definition of intended use, performance expectations, oversight, and accountability [16].

5.3 Practical controls that determine success

Across implementations, durable AI value correlates with:

- **Data quality programs** (standardization, reduplication, feedback loops)
- **Model risk management** (intended use, validation, drift monitoring)
- **Human-in-the-loop designs** (confidence thresholds, escalation rules)
- **Audit readiness** (versioning, decision logs, traceability to source text)
- **Privacy/security** (especially for cloud NLP and multi-tenant tooling)

Table 1. Evolution of pharmacovigilance from manual to AI-enabled operations

Era	Core PV operating model	Typical data	Key limitations	Enablers of next shift
Manual / paper-centric	Human extraction + manual coding + narrative review	Paper forms, faxes, PDFs	Latency, inconsistency, scaling bottlenecks	Digitization, centralized databases
Electronic / standardized	Structured ICSR exchange; standard terminologies	E2B(R3) messages; MedDRA-coded events	Still labor intensive; unstructured narrative underused	NLP, better data pipelines
Augmented PV	AI supports extraction, coding, prioritization	Hybrid structured + free text	Model drift, bias, explain ability gaps	Governance frameworks, continuous monitoring
AI-enabled surveillance	Near-real-time triage, automated coding at scale, hybrid signal intelligence	Multi-source + active surveillance	Integration complexity; accountability	Mature MRM, robust QA, audit ability

Key standardization foundations include ICH E2B (R3) for ICSR exchange and MedDRA for standardized coding.

Table 2. AI methods mapped to PV tasks, validation metrics, and risks

PV task	AI approach	Typical metrics	Expected operational benefit	Primary risks / controls
Valid case identification / triage	Text classification; rules + ML	Sensitivity/recall (must be high), PPV	Faster routing; reduced backlog	False negatives; use conservative thresholds + QA sampling
Entity extraction from narratives	NER transformers; negation handling	F1, span accuracy	Less manual transcription	Language drift; maintain labeled sets, monitor drift
Drug & event coding	Dictionary matching + ML ranking	Top-1/Top-k accuracy	Throughput; consistency	Over coding/under coding; "code when confident" policy
Duplicate detection	Similarity models; embedding	Precision at high recall	Reduced duplicate noise	Incorrect merges; require reviewer confirmation
Signal prioritization	Hybrid ML + disproportionality	AUROC + calibration	Focus reviewer time	Black-box risk; require explainable features & dashboards
Social media screening	NLP classifiers; ADR extraction	Precision (high), triage yield	Early weak signals	Noise; triangulate with trusted sources

Examples of automation in PV include WHO Drug Koda evaluations for automated drug coding and deep-learning approaches for case processing augmentation.

Table 3. Governance checklist for implementing AI in pharmacovigilance (submission-grade)

Domain	Minimum control set	Evidence artifacts (audit-ready)
Intended use & scope	Define task boundaries; exclusions; escalation	Use-case SOP; model card; risk assessment
Data governance	Provenance, labeling protocol, access controls	Data dictionary; lineage; QC logs
Validation & performance	Pre-deployment validation; subgroup checks	Validation report; bias assessment; test sets
Human oversight	Confidence thresholds; override workflow	Review logs; sampling plan; escalation records

Change control	Versioning; release approvals; rollback	Change requests; release notes; approvals
Continuous monitoring	Drift metrics; periodic re-validation	Monitoring dashboards; periodic QA reports
Regulatory alignment	Map to E2B(R3)/GVP obligations	Compliance mapping; inspection-ready pack

CIOMS WG XIV highlights the need for oversight, trustworthiness, and accountability when deploying AI in PV.

6. DISCUSSION

6.1 What is truly "reinvented" in AI-enabled PV?

The most meaningful reinvention is not replacing medical judgment but restructuring work so that experts spend time on interpretation rather than transcription. Across organizations, the highest-ROI applications tend to be:

- automated coding (products/events),
- extraction of structured fields from narratives,
- duplicate detection support,
- Prioritization of queues and candidate signals.

These are high-volume; repetitive tasks with measurable quality and efficiency outcomes.

6.2 Why many AI pilots fail to scale

Common causes include:

- **Weak training labels** (inconsistent "ground truth" from heterogeneous processors)
- **Over fitting to one vendor's narratives or one region/language**
- **Lack of end-to-end metrics** (time saved in the workbench, rework rate, compliance impact)
- **Insufficient governance** (no drift monitoring, no change control, poor audit trails)

6.3 The emerging role of LLMs

Large language models can summarize narratives, draft case narratives, assist medical reviewers, and support literature triage. However, PV requires conservative deployment patterns: retrieval-augmented workflows, strict source traceability, and prevention of hallucinations. In practice, LLMs are best treated as **assistive** layers with mandatory citation to source documents and constrained outputs.

6.4 Equity, bias, and global generalizability

PV is inherently global. AI systems trained on reports from one geography or reporter type may underperform elsewhere. Subgroup evaluation (age, sex, language, and region) and continuous monitoring are essential to prevent systematic under-detection in under-represented populations.

7. CONCLUSION

Pharmacovigilance has moved from manual, document-centric operations to standardized electronic reporting and now to AI-enabled surveillance. The enabling infrastructure—ICSR standards (ICH E2B (R3)), global databases (FAERS, EudraVigilance, VigiBase), and controlled terminologies (MedDRA)—makes it feasible to deploy AI at scale, particularly for extraction, coding, and prioritization. The dominant implementation lesson is that **augmentation plus governance** outperforms "full automation": organizations realize the most sustainable gains when AI is embedded into inspection-ready workflows with human oversight, performance monitoring, and strong data governance. Future progress will depend less on novel algorithms and more on reliable operating models: transparent validation, drift detection, auditability, and aligned regulatory expectations.

REFERENCES

1. International Council for Harmonization (ICH). E2B (R3) Individual Case Safety Report (ICSR) Specification and related files. ICH website.
2. U.S. Food and Drug Administration (FDA). FDA Adverse Event Reporting System (FAERS) Public Dashboard. Updated Dec 7, 2023.
3. European Medicines Agency (EMA). EudraVigilance: system for managing and analyzing suspected adverse reactions in the EEA.
4. Uppsala Monitoring Centre (UMC). About VigiBase: WHO global database of individual case safety reports? Updated Oct 17, 2025.
5. FDA. Sentinel Initiative: national electronic system for proactive monitoring. Updated Mar 8, 2024.

6. Platt R, et al. The FDA Sentinel Initiative—an evolving national resource. *N Engl J Med*. 2018. doi:10.1056/NEJMp1809643.
7. CIOMS Working Group XIV. *Artificial intelligence in pharmacovigilance* (draft report for public consultation). May 1, 2025.
8. Nagar A. Artificial intelligence in pharmacovigilance: advancing drug safety monitoring (review). *[PMCID article]*. 2025.
9. Algarve RC, et al. Artificial intelligence in pharmacovigilance: a narrative review. *[PMCID article]*. 2025.
10. EMA. Guideline on good pharmacovigilance practices (GVP) Module VI – Collection, management and submission of reports of suspected adverse reactions to medicinal products (Rev. 2). Legal effective date Nov 22, 2017MedDRA. Medical Dictionary for Regulatory Activities—standardized terminology for regulatory information sharing. MedDRA/ICH resources.
11. Meldau EL, et al. Automated drug coding using artificial intelligence: an evaluation of WHODrug Koda on adverse event reports. *Drug Safe*. 2022; 45(5):549–561. doi:10.1007/s40264-022-01162-7.
12. Abate Marco D, Perera S, Bao SH, et al. Training augmented intelligent capabilities for pharmacovigilance: applying deep-learning approaches to individual case safety report processing. *Pharm Med*. 2018; 32(6):391–401. doi:10.1007/s40290-018-0251-9.
13. Nikfarjam A, et al. Pharmacovigilance from social media: mining adverse drug reaction mentions using sequence labeling with word embedding cluster features. *J Am Med Inform Assoc*. 2015.
14. De Rosa M, et al. Pharmacovigilance in the era of social media. *Future Genre Compute Syst*. 2021.
15. CIOMS. Working Group XIV: Artificial Intelligence in Pharmacovigilance—objectives and principles.
16. WHO. Pharmacovigilance (PVG archive): WHO Programmed for International Drug Monitoring and VigiBase overview.
17. FDA. OpenFDA FAERS dataset overview.
18. Postigo R, et al. EudraVigilance medicines safety database: publicly accessible data and regulatory transparency. *[PMCID article]*. 2018.
19. Shamim MA, et al. Artificial intelligence and big data for pharmacovigilance (review). *[Science Direct article]*. 2024.