

# AI in Pharmacovigilance: Automating Adverse Event Detection

By Adrien Laurent, CEO at IntuitionLabs • 2/3/2026 • 40 min read

pharmacovigilance, artificial intelligence, adverse event detection, gvp guidelines, natural language processing

drug safety, machine learning, regulatory compliance, signal detection, icsr processing



# Executive Summary

Pharmacovigilance (PV) – the science of detecting, assessing, and preventing adverse effects of medicines – has long relied on laborious manual processes. In recent years, *artificial intelligence (AI)* and advanced data-processing techniques have begun to transform PV by automating the detection of adverse drug events (AEs) across vast datasets. This shift is driven by an explosion in safety data: for example, one major analysis notes that some marketing authorization holders (MAHs) process **over one million individual case safety reports (ICSRs)** per year (<sup>[1]</sup> pmc.ncbi.nlm.nih.gov). Traditional case processing can consume up to **two-thirds of a company's PV budget** (<sup>[2]</sup> pmc.ncbi.nlm.nih.gov), highlighting the need for efficiency gains.

AI offers powerful tools – including **machine learning (ML)**, **natural language processing (NLP)**, **deep neural networks**, and **robotic process automation (RPA)** – that can extract relevant information from unstructured texts, mine signals from structured databases, and support decision-making. Case studies and pilot projects demonstrate the promise of these technologies. Notably, a 2018 Pfizer pilot successfully used AI to extract key report elements (CAS, drug, patient, reporter) from source documents (<sup>[3]</sup> pmc.ncbi.nlm.nih.gov), and a 2022 French initiative reported an AI system achieving area-under-ROC-curve (AUC) ≈0.97 for ADR identification and high scores for seriousness coding in patient reports (<sup>[4]</sup> pmc.ncbi.nlm.nih.gov). These tools have even been deployed in practice (e.g. by French regulators for COVID-19 vaccine monitoring (<sup>[5]</sup> pmc.ncbi.nlm.nih.gov)).

While AI can **accelerate and enhance safety monitoring**, its adoption in PV must adhere to strict regulatory standards. In Europe, Good Pharmacovigilance Practices (GVP) modules provide the legal framework for safety reporting, quality systems, and signal management ([www.ema.europa.eu](http://www.ema.europa.eu)) ([www.ema.europa.eu](http://www.ema.europa.eu)). Under GVP (and analogous FDA guidelines), validated IT systems are required; indeed, PV regulations mandate that pharmacovigilance systems be “fit for purpose” and **qualified/validated** (<sup>[6]</sup> [www.biologit.com](http://www.biologit.com)). AI tools used in PV must therefore be rigorously tested, transparent, and integrated with human oversight. This report examines how AI is being used to automate adverse event detection *within* the GVP framework, balancing potential benefits against challenges such as data quality, algorithmic bias, and interpretability.

Key findings include:

- **Data Volume and Need for Automation.** Worldwide **safety databases** have grown exponentially. For instance, the US FDA's FAERS database now contains over 24 million ICSRs with ~2 million new reports annually (<sup>[7]</sup> [intuitionlabs.ai](http://intuitionlabs.ai)). The exponential growth of PV data – driven by regulatory requirements and digital reporting – has made manual workflows increasingly strained (<sup>[8]</sup> [intuitionlabs.ai](http://intuitionlabs.ai)) (<sup>[9]</sup> pmc.ncbi.nlm.nih.gov). AI is seen as essential to “separate needles from haystack” and to maintain timely monitoring at scale (<sup>[8]</sup> [intuitionlabs.ai](http://intuitionlabs.ai)) (<sup>[9]</sup> pmc.ncbi.nlm.nih.gov).
- **AI Methods and Applications.** A broad spectrum of AI techniques are under study or in use. **NLP** (for parsing free text in ICSRs, literature, EHRs) and **ML classifiers** (for case triage, coding, signal detection) are well-established. Emerging applications include **knowledge graphs** linking drug-event networks and modern **large language models** for summarizing complex narratives. Common tasks include automated coding of medical terms (e.g. using MedDRA) (<sup>[4]</sup> pmc.ncbi.nlm.nih.gov), deduplication of ICSRs (e.g. UMC's vigiMatch system) (<sup>[10]</sup> [uppsalareports.org](http://uppsalareports.org)), and preliminary assessment of seriousness or causality (<sup>[11]</sup> pmc.ncbi.nlm.nih.gov).
- **Performance and Case Studies.** AI-driven PV tools have shown high accuracy in controlled studies. One Ferrer et al. study reported F1-scores up to ~0.74 in extracting all key fields from source documents (<sup>[3]</sup> pmc.ncbi.nlm.nih.gov). An IBM/Celgene collaboration achieved 83–93% accuracy in automated identification of seriousness across various report types (<sup>[12]</sup> pmc.ncbi.nlm.nih.gov). In France, AI models reached AUC=0.97 for ADR identification (and ~0.85 for seriousness) in patient-submitted reports (<sup>[4]</sup> pmc.ncbi.nlm.nih.gov). Another review summarizing real-world use noted that Bayesian-network algorithms reduced causality assessment times from “days to hours” at a PV center (<sup>[13]</sup> pmc.ncbi.nlm.nih.gov).

- **Regulatory Compliance under GVP.** AI tools must comply with PV regulations (GVP in EU, FDA guidelines in US). For example, GVP Module VI and its addenda govern ICSR collection and data submission, including duplicate handling ([www.ema.europa.eu](http://www.ema.europa.eu)). Module I requires quality systems that encompass all PV software (<sup>[6]</sup> [www.biologit.com](http://www.biologit.com)). Guidance (CIOMS, ICMRA, FDA/ICH guidelines) emphasizes rigorous validation, transparency, and human oversight for PV AI systems (<sup>[6]</sup> [www.biologit.com](http://www.biologit.com)) (<sup>[14]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)).
- **Challenges and Risks.** Critical challenges include *data quality*, *algorithmic bias*, and *explainability*. AI models trained on historical ICSRs may inherit reporting biases, and may underperform on rare events. Black-box algorithms raise concerns for auditability, especially under regulations that require understanding a system's decision logic. Ethical issues like patient privacy (e.g. personal data masking requirements in GVP Module VI Addendum II) also constrain AI usage. Several reviewers note that "AI should augment, not replace" PV experts (<sup>[15]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)) (<sup>[14]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)); building trust and clarity are essential.
- **Future Implications.** Looking ahead, generative AI and advanced analytics could further revolutionize PV – for instance, by auto-summarizing new research or extracting signals from **real-world data**. EU and global regulators are actively preparing guidelines (e.g. CIOMS WG XIV has released a 2025 consensus report on AI in PV) to ensure AI's safe **integration**. As one expert notes, the upcoming years will "define the foundation for trustworthy, explainable, and regulatory-compliant AI in pharmacovigilance" ([cioms.ch](http://cioms.ch)) (<sup>[14]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)).

In summary, automating adverse event detection with AI under GVP promises significant efficiency and safety benefits, but must be pursued carefully. This report provides an in-depth review of technologies, evidence, case examples, regulatory context, and future directions to guide stakeholders in leveraging AI for pharmacovigilance.

## Introduction

Pharmacovigilance (PV) is a regulatory-mandated, post-marketing process to ensure medicine safety. It involves gathering and analyzing data on **suspected adverse drug reactions (ADRs)** or any untoward effects following medication use. The objective is to **detect safety signals early**, assess risk versus benefit, and take action if needed. Historically, PV relied on spontaneous case reports submitted by healthcare professionals or patients, which are compiled into databases such as the WHO's VigiBase or the FDA's FAERS. These datasets feed traditional **signal detection** methods (e.g. disproportionality analyses) and inform regulatory decision-making.

The **volume and complexity of safety data** have grown tremendously. Digital reporting systems, electronic health records (EHRs), patient support programs, and even social media have emerged as new sources of information. One analysis cites that individual marketing authorization holders may handle **over one million safety-related transactions** per year, including ICSRs, medication error reports, and product quality complaints (<sup>[1]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)). The FDA's FAERS database alone contains billions of data fields – at least 24 million ICSRs to date, with ~2 million new entries per year (<sup>[7]</sup> [intuitionlabs.ai](http://intuitionlabs.ai)). This "data deluge" strains traditional workflows. Multiple observers note that manually triaging these reports is increasingly impractical, likening it to finding "needles in a haystack" (<sup>[8]</sup> [intuitionlabs.ai](http://intuitionlabs.ai)) (<sup>[9]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)).

Good Pharmacovigilance Practices (GVP) provide the regulatory framework for managing this data. In the European Union, the EMA's GVP modules (I–XVI) define how marketing holders must collect, analyze, and report safety information ([www.ema.europa.eu](http://www.ema.europa.eu)) ([www.ema.europa.eu](http://www.ema.europa.eu)). For example, **GVP Module VI** covers the *collection and submission of ICSRs* (including quality checks and duplicate management ([www.ema.europa.eu](http://www.ema.europa.eu))), while **Module IX** addresses *signal management* including statistical detection ([www.ema.europa.eu](http://www.ema.europa.eu)). GVP and analogous FDA guidelines (e.g. 21 CFR Part 314.80) require that any PV system be robust and validated. Per GVP, "IT systems used in pharmacovigilance should be fit for purpose, and subject to appropriate checks, qualification and/or validation activities to prove their suitability" (<sup>[6]</sup> [www.biologit.com](http://www.biologit.com)). In other words, whether processing is done by humans or AI, the integrity of the PV system must be assured.

Given these demands, the advent of AI and machine learning offers a compelling solution. AI encompasses a suite of computational approaches (from rule-based engines to deep learning) that can learn patterns in data and automate reasoning (<sup>[16]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)) (<sup>[17]</sup> [uppsalareports.org](http://uppsalareports.org)). In PV, AI tools aim to **augment human expertise** by

performing repetitive data tasks faster and spotting subtle correlations. For instance, natural language processing (NLP) can extract structured information from narrative case reports or scientific articles (<sup>[18]</sup> [pmc.ncbi.nlm.nih.gov](https://pmc.ncbi.nlm.nih.gov)) (<sup>[19]</sup> [pmc.ncbi.nlm.nih.gov](https://pmc.ncbi.nlm.nih.gov)); classification algorithms can flag likely duplicate cases (<sup>[10]</sup> [uppsalareports.org](https://uppsalareports.org)), and predictive models can prioritize safety signals before manual review. As one expert commentary observes, AI promises to transform safety monitoring by “streamlin(ing) signal detection and surveillance” and automating AED (adverse event detection) to expedite risk identification (<sup>[20]</sup> [pmc.ncbi.nlm.nih.gov](https://pmc.ncbi.nlm.nih.gov)).

However, these technologies must be evaluated within the PV regulatory environment. Automation cannot compromise compliance: any AI system must still meet GVP requirements for accuracy, traceability, and oversight. This means rigorous model validation, documentation of decision rules, and integration of human review (“hybrid intelligence”) to handle uncertainty. The Council for International Organizations of Medical Sciences (CIOMS) notes the need to treat PV AI agents “like medicinal products” – with clearly defined scope, capabilities, and limitations ([cioms.ch](https://cioms.ch)). In short, AI-for-PV is a “high stakes” undertaking that requires careful governance (<sup>[21]</sup> [www.biologit.com](https://www.biologit.com)).

This report examines the intersection of AI and pharmacovigilance, focusing on automating adverse event detection within the framework of Good Pharmacovigilance Practices. We review the state of the art in AI methods for PV, provide data-driven examples and case studies, discuss regulatory and technical considerations, and explore future directions. Emphasis is placed on evidence-based outcomes (with numerous examples of study results and expert analyses) and on compliance with established PV guidelines.

## Pharmacovigilance and Regulatory Framework

### Good Pharmacovigilance Practices (GVP) and Global Guidelines

Good Pharmacovigilance Practices (GVP) in the EU, instituted by the 2010 Pharmacovigilance Legislation, consist of modules that describe how to set up and operate PV systems ([www.ema.europa.eu](https://www.ema.europa.eu)). GVP Module I mandates that MAHs maintain a *pharmacovigilance system master file* and an accompanying **quality system** to ensure data integrity ([www.ema.europa.eu](https://www.ema.europa.eu)). Crucially, GVP requires that all PV processes (including electronic components) be documented, validated, and audited. As noted in guidance:

“IT systems used in pharmacovigilance should be fit for purpose, and subject to appropriate checks, qualification and/or validation activities to prove their suitability.” (<sup>[6]</sup> [www.biologit.com](https://www.biologit.com))

Module VI focuses on *case processing* — collection, management and submission of ICSRs. It details content requirements (patient, reporter, drug, reaction, etc.) and highlights procedures for *duplicates* and *seriousness evaluation* ([www.ema.europa.eu](https://www.ema.europa.eu)). Addenda to Module VI cover specific topics like duplicate case management (Addendum I) and personal data masking (Addendum II in 2025) ([www.ema.europa.eu](https://www.ema.europa.eu)). Module IX addresses *signal management*, including methodological aspects of detecting signals from spontaneous reports ([www.ema.europa.eu](https://www.ema.europa.eu)). Entities must follow standardized case definitions (ICH E2B(R3) format for ICSRs) and timelines per GVP rules.

Beyond Europe, regulatory authorities worldwide are formulating AI-relevant guidance. The FDA has proposed a Total Product LifeCycle approach for AI/ML-based software as a medical device (SaMD) (<sup>[22]</sup> [www.biologit.com](https://www.biologit.com)), and issued Good Machine Learning Practices (GMLP) emphasizing transparency, validation, and post-market monitoring (<sup>[23]</sup> [www.biologit.com](https://www.biologit.com)) (<sup>[24]</sup> [www.biologit.com](https://www.biologit.com)). In 2024, CIOMS WG XIV released a *Consensus Report on AI in Pharmacovigilance*, explicitly addressing AI use cases, risk management, and the need for clear indications and limitations ([cioms.ch](https://cioms.ch)) ([cioms.ch](https://cioms.ch)). The International Coalition of Medicines Regulatory Authorities (ICMRA) likewise recommends multi-disciplinary oversight and continuous model performance monitoring in PV AI systems (<sup>[24]</sup> [www.biologit.com](https://www.biologit.com)).

Importantly, these guidelines consistently underscore the principle of **human-AI collaboration**. For example, a recent Indian review notes that AI should “augment, not substitute, human experts” in PV (<sup>[15]</sup> [pmc.ncbi.nlm.nih.gov](https://PMC.ncbi.nlm.nih.gov)). EMA and FDA emphasize *interpretability*, user training, error rates, and customer feedback loops. In essence, automating AE detection must occur *under the auspices of established quality management*, with fallback to human review as needed. The shift is not to eliminate PV professionals, but to free them from routine tasks so they can focus on complex signal assessment.

## Data Sources and Systems

Pharmacovigilance draws on multiple data streams. The core is **Individual Case Safety Reports (ICSRs)** submitted to national and international databases (e.g., FDA FAERS, EMA EudraVigilance, WHO VigiBase). Additional sources include **Periodic Safety Update Reports (PSURs)** from sponsors, post-authorization safety studies, literature case reports, EHRs and claims data (for pharmacoepidemiology), and increasingly **digital data** (e.g. social media, forums, mobile apps). All these feed into PV systems that generate signals. For example, WHO's VigiBase contains over 35 million ICSRs (as of 2026) from 150+ countries, spanning decades (<sup>[25]</sup> [pmc.ncbi.nlm.nih.gov](https://PMC.ncbi.nlm.nih.gov)). Managing and mining this variety of data is a central PV challenge.

With the EU's new digital health regulations and EudraVigilance since 2017, reporting has become more structured and voluminous. Meanwhile, patient engagement portals (like France's [signalement.social-sante.gouv.fr](https://signalement.social-sante.gouv.fr)) increased public reporting. The bottom line is that PV systems now accumulate disparate data, much of it **unstructured text** (narrative fields, literature). AI techniques, especially NLP, are therefore highly applicable to PV tasks.

## Challenges of Traditional Pharmacovigilance

Traditional PV workflows involve **case intake, data entry, medical review, coding and reporting, and signal detection**. Each ICSR typically goes through multiple manual steps. Case processing (the series of tasks around an ICSR) is by far the largest PV expense. Several sources report that processing ICSRs can consume **50–70% of a PV budget** (<sup>[2]</sup> [pmc.ncbi.nlm.nih.gov](https://PMC.ncbi.nlm.nih.gov)) (<sup>[26]</sup> [intuitionlabs.ai](https://intuitionlabs.ai)). Considering that with pharmacovigilance obligations expanding, even large safety teams can be overwhelmed, delays may occur. Critical decisions (e.g. causality assessment, seriousness determination) then rely on human judgment that can be influenced by fatigue or bias.

A wealth of studies echoes the “data tsunami” narrative. For instance, in India, case reports to the national PV program (PvPI) have grown *exponentially*, mirroring WHO's global increase (<sup>[9]</sup> [pmc.ncbi.nlm.nih.gov](https://PMC.ncbi.nlm.nih.gov)). The manual effort to sift through these is immense: “*nearly 1 in 5 [PV teams] rely on manual or outdated methods*”, and over 65% describe their approach as “*largely reactive*” (<sup>[27]</sup> [intuitionlabs.ai](https://intuitionlabs.ai)). This creates a bottleneck for real-time safety surveillance. Behind the scenes, large pharmaceutical firms have similarly documented burdens: one analysis noted that “case processing activities constitute a significant portion of PV resource use, up to two-thirds” (<sup>[2]</sup> [pmc.ncbi.nlm.nih.gov](https://PMC.ncbi.nlm.nih.gov)).

These pressures are not solely internal. Public and regulatory expectations demand faster detection of safety signals. In the context of COVID-19 vaccination, for example, the unprecedented volume of reports “catalyzed a digital transformation”, with regulators seeking novel tech solutions ([cioms.ch](https://cioms.ch)). Governments are demanding real-time pharmacovigilance to identify even rare events swiftly. Therefore, both patient safety and compliance are driving adoption of advanced IT.

## AI and Machine Learning in Pharmacovigilance

## Definitions and Scope

**Artificial Intelligence (AI)** is broadly defined as the simulation of human intelligence processes by machines. Within AI, **Machine Learning (ML)** refers to algorithms that improve performance on a task with experience (data). For PV, ML encompasses techniques from logistic regression, decision trees, Bayesian models, to deep learning neural networks. A useful taxonomy is given by expert Aronson: “*AI includes any application of machine learning and natural language processing, as well as expert systems explicitly programmed to perform specific tasks*” (<sup>[17]</sup> [uppsalareports.org](http://uppsalareports.org)).

Key AI modalities in PV include:

- **Natural Language Processing (NLP):** Enables understanding human-language text. NLP can extract structured data (names of drugs, dosage, reactions) from free-text fields in ICSRs or from unstructured sources like literature and social posts (<sup>[18]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)) (<sup>[28]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)). For example, entity-recognition models can parse narratives to populate the standard PV database fields.
- **Neural Networks / Deep Learning:** Complex models (e.g. convolutional or recurrent NNs) that can learn hierarchical features. Useful for classification tasks (e.g. categorizing an AE seriousness) and for processing sequential data like text or audio transcripts.
- **Knowledge Graphs / Network Analysis:** Graph-based AI can link drugs, targets, phenotypes, and patient demographics to infer hidden relationships. Network methods were used to study complex ADRs (e.g., impulse control disorders) (<sup>[29]</sup> [uppsalareports.org](http://uppsalareports.org)).
- **Robotic Process Automation (RPA):** Not AI per se, but often associated. RPA uses scripts/bots to automate rule-based tasks (e.g. copying data from emails into PV databases). When combined with AI modules, chores can be mostly automated.
- **Predictive Analytics:** Supervised models trained to predict ADRs based on existing data, or unsupervised models (clustering, anomaly detection) to flag unusual patterns.

The goals of AI in PV emphasize **efficiency, quality, and new capabilities** (<sup>[30]</sup> [uppsalareports.org](http://uppsalareports.org)). Efficiency means doing tasks faster (e.g. automatic triage), quality means more accurate or complete data, and capability means unlocking insights not feasible manually (发现 hidden safety signals). Fundamentally, PV AI is regarded as *intelligence augmentation* – the machine assists experts, who retain ultimate judgment (<sup>[31]</sup> [uppsalareports.org](http://uppsalareports.org)).

Table 1 below summarizes common PV tasks and corresponding AI techniques:

| PV Task / Process                            | Typical AI/Machine Learning Solution                              | Benefits/Examples  |
|--|---|--|
| Case Intake & Triaging (ICSR validity)       | NLP extraction, supervised classification, RPA routing            | Quickly verify report completeness; flag valid vs invalid cases. Example: Automated ICSR validation using ML saved manpower ( <sup>[3]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ).   |
| Coding of Drugs and Reactions (MedDRA, etc.) | Deep learning NLP, Transformer models                             | Map free-text terms to standardized codes. Achieved high performance (AUROC=0.97*) in coding patient reports ( <sup>[4]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ).  |
| Duplicate Case Detection                     | Graph algorithms, similarity measures (vigiMatch)                 | Identify duplicate ICSRs to avoid double counting. UMC's ML-based deduplication improved accuracy at scale ( <sup>[10]</sup> <a href="http://uppsalareports.org">uppsalareports.org</a> ).   |
| Seriousness & Expectedness Assessment        | Neural network classifiers, rule-enabled ensemble                 | Classify AEs as serious (hospitalization, death, etc.). Deep models achieved >80% accuracy on broad case sets ( <sup>[12]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ).  |
| Signal Detection (Spontaneous Reports)       | Disproportionality (with ML), Bayesian methods, cluster detection | Detect elevated drug-event associations. AI can augment traditional ROR/PRR by pattern mining ( <sup>[32]</sup> <a href="http://uppsalareports.org">uppsalareports.org</a> ) ( <sup>[20]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ). |
| Literature & Social Media Monitoring         | Large-scale NLP/ML (e.g. transformer models)                      | Scan journals, clinical trial registries, or tweets/posts for new ADR mentions. Improves completeness of data ( <sup>[28]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ).  |
| Aggregate Data Analysis (EHR, Registries)    | Predictive models, multivariate analysis                          | Identify ADR trends in real-world data. E.g. automated scanning of EHRs to detect rare AEs (ongoing research).   |

Table 1. Examples of AI/ML applications in pharmacovigilance and their benefits (ICSR = Individual Case Safety Report). Many of these have been evaluated or implemented in research trials and pilot systems (<sup>[3]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)) (<sup>[4]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)) (<sup>[32]</sup> [uppsalareports.org](http://uppsalareports.org)) (<sup>[12]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)).

## Machine Learning Models and Data

AI models require training data. In PV contexts, labeled datasets are scarce due to privacy and quality issues.

Nevertheless, several strategies help build training corpora:

- **Expert-annotated corpora:** For example, Pfizer's 2018 pilot trained NLP models on a set of ICSRs annotated by safety experts (<sup>[3]</sup> pmc.ncbi.nlm.nih.gov). However, manual annotation of thousands of reports is cost-prohibitive.
- **Database-derived labels:** Some pilots (e.g. the same Pfizer study) used fields in the company's safety database as *surrogates* for annotation, leveraging pre-coded information (<sup>[33]</sup> pmc.ncbi.nlm.nih.gov).
- **Crowdsourcing / Public Data:** A notable French study collected 11,633 patient reports (with known outcomes) from the national PV database to train ADR identification models (<sup>[34]</sup> pmc.ncbi.nlm.nih.gov). This kind of consolidated dataset (like national PV centers' archives) is ideal for training and internal validation.
- **Transfer Learning:** Recent approaches apply pretrained language models (BERT variants, etc.) for PV text tasks. For instance, the French team fine-tuned CamemBERT for French-language report coding (<sup>[4]</sup> pmc.ncbi.nlm.nih.gov).

The **data challenges** in PV are substantial. Reports often lack key fields or contain jargon and abbreviations. NLP systems must handle multiple languages (especially in global databases). Moreover, spontaneous-report data have biases (e.g. under-reporting of common mild AEs) that AI can inadvertently amplify. Consequently, model performance is carefully validated on held-out data. A safety threshold is often set: if ML accuracy/F1 falls below a pre-defined cutoff (e.g. 75%), human review is required (<sup>[35]</sup> pmc.ncbi.nlm.nih.gov).

Despite challenges, results have been encouraging. In one case, several vendors' AI systems were able to extract all critical fields and identify valid AE cases with average composite F1 ≈0.7 (<sup>[3]</sup> pmc.ncbi.nlm.nih.gov). Another characterization: "AI-based technology is viable to support extraction from AE source documents and evaluation of case validity" (<sup>[3]</sup> pmc.ncbi.nlm.nih.gov). Importantly, robust validation frameworks (exclusive test sets, cross-validation, post-deployment monitoring) are emphasized at every step (<sup>[35]</sup> pmc.ncbi.nlm.nih.gov).

## Applications of AI to Case Processing

A major area of application is automating **case intake and processing** of ICSRs, particularly extracting information and completing database fields. Case processing comprises tasks such as verifying report validity, identifying the suspect drug and AE terms, assigning MedDRA codes, determining event seriousness, and populating other regulatory-required data.

One of the first industry pilots (Pfizer, 2018) demonstrated feasibility of AI for this workflow. The team tested three commercial NLP/ML systems on a sample of ICSRs, using the company's existing database records as ground truth instead of hand annotation (<sup>[3]</sup> pmc.ncbi.nlm.nih.gov). The AI systems achieved overall extraction accuracies (F1) in the 0.52–0.74 range, with the best systems outperforming an internal benchmark (<sup>[3]</sup> pmc.ncbi.nlm.nih.gov). Specifically, "AI-based technology" successfully extracted core elements (AE, causal drug, patient info, reporter) and flagged case validity in compliance with regulatory definition (<sup>[3]</sup> pmc.ncbi.nlm.nih.gov). These findings affirmed that careful machine-learning pipelines could match human performance on many case fields, suggesting large potential labor savings.

**Coding and terminology.** Automated coding of free-text to standard dictionaries (e.g. MedDRA for medications and AEs) is a high-impact area. Traditional coding is tedious and error-prone. State-of-art ML pipelines (combining text embeddings and classifiers) have been validated. In the 2022 French study (<sup>[4]</sup> pmc.ncbi.nlm.nih.gov), the top-ranking model (TF-IDF + LightGBM) achieved **AUC ≈0.97** for identifying patient-reported ADRs, and *the same sensitivity* as human coders. The tool was good enough that French regulators **deployed it since Jan 2021** to speed up analyses, particularly during COVID-19 vaccine monitoring (<sup>[5]</sup> pmc.ncbi.nlm.nih.gov). This real-world use-case shows that when performance is high and domain coverage is adequate, regulators are willing to incorporate AI into official pipelines.

**Seriousness and triage.** Deciding whether an event is “serious” (death, hospitalization, etc.) is crucial in PV as it drives reporting timelines. AI models have been trained for this task across report types. In one multi-institutional study (<sup>[12]</sup> pmc.ncbi.nlm.nih.gov), a deep-learning classifier reached 83–93% accuracy in automatically flagging report-level seriousness on spontaneous, solicited, and literature reports. The conclusion was that “a neural network approach can provide an accurate and scalable solution” to assist human workers (<sup>[36]</sup> pmc.ncbi.nlm.nih.gov). Another industry effort (Novartis in collaboration with IBM) achieved similarly high F1-scores (~0.86) for event seriousness classification using recurrent neural networks (<sup>[12]</sup> pmc.ncbi.nlm.nih.gov) (<sup>[11]</sup> pmc.ncbi.nlm.nih.gov). These AI scorings can triage cases faster, ensuring no serious case is overlooked due to backlog.

**Duplicate detection.** Identifying duplicate ICSRs for the same patient and event is important to avoid inflating signal strength. AI methods here include probabilistic record linkage and advanced pattern analysis. The Uppsala Monitoring Centre (WHO-UMC) reports that **vigiMatch**, an ML-driven duplicate-finding system, can process around **50 million report-pairs per second**, far beyond human capability (<sup>[37]</sup> uppsalareports.org). In practice, even simple AI pattern-matching algorithms have significantly reduced duplicate reporting errors. This is an early and illustrative success of PV automation.

**Outcome.** Overall, AI applications in case processing aim to yield “high quality safety data in the correct format, in context, more quickly, and with less manual effort” (<sup>[38]</sup> pmc.ncbi.nlm.nih.gov). Table 2 summarizes select examples of AI applied to PV case tasks:

| Study / Implementation  | Task   | Data / Setting                    | AI Technique(s)                      | Performance / Status   |
|---|--|-----------------------------------|--------------------------------------|--|
| Schmid et al., 2018 (Pfizer pilot) ( <sup>[3]</sup> pmc.ncbi.nlm.nih.gov) | AE case intake, ICSR processing                          | Internal ICSRs (multiple vendors) | NLP + ML (vendor RPA tools)          | F1 = 0.52–0.74 across fields; top system F1 = 0.74 for all fields; viability proven. |
| Routray et al., 2019 ( <sup>[12]</sup> pmc.ncbi.nlm.nih.gov)              | Classify AE seriousness (binary & categorization)        | Celgene ICSRs                     | Recurrent Neural Network             | Accuracy 83–93% (varying by report type); scalable support for human review.         |
| Martin et al., 2022 (France) ( <sup>[4]</sup> pmc.ncbi.nlm.nih.gov)       | Automate ADR identification & coding in patient reports  | 11,633 French ADR reports         | TF-IDF + LightGBM; XLM (transformer) | AUC = 0.97 for ADR ID, F=0.80; tool deployed by health authorities (COVID PV).       |
| Fusaroli et al., 2024 (UMC) ( <sup>[10]</sup> uppsalareports.org)         | ICSRs duplicate detection                                | WHO Global ICSRs                  | Custom ML deduplication (vigiMatch)  | Processes ~50M report pairs/sec; early AI scaling proven.                            |
| Algarvio et al., 2025 ( <sup>[13]</sup> pmc.ncbi.nlm.nih.gov)             | Automate causality/seriousness assessment (Bayesian net) | Regional PV center pilot          | Expert-defined Bayesian Network      | Reduced causality review time from days to hours; improved consistency.              |

Table 2. Examples of AI-driven projects in pharmacovigilance case processing and signal management (AE = adverse event; ADR = adverse drug reaction; ICSR = Individual Case Safety Report).

Each of these projects demonstrates that AI can handle key PV tasks at or above human-expert performance. The success often depends on *integrating domain expertise* (e.g. human-in-the-loop training or rules) and rigorous validation. As a result, many PV professionals anticipate steady adoption *within the GVP framework*: completing lower-level tasks automatically while experts focus on judgment-intensive analysis (<sup>[3]</sup> pmc.ncbi.nlm.nih.gov) (<sup>[14]</sup> pmc.ncbi.nlm.nih.gov).

## AI for Signal Detection and Safety Analysis

While case processing is about handling individual reports, **signal detection** looks for *statistical associations or trends* indicating new safety issues. Traditionally, this relies on disproportionality analyses (e.g. reporting odds ratios) in spontaneous-report databases (<sup>[39]</sup> pmc.ncbi.nlm.nih.gov). AI and advanced analytics can enhance this by uncovering subtler patterns.

**Disproportionality with ML enhancements.** AI can automate the scanning of large databases for unusual drug–event pairs. Some systems use *machine-learning enhanced statistics*: for instance, training a classifier to predict whether a

given drug–event pair is likely a true signal using historical data as labels. Others incorporate Bayesian methods or ensemble approaches. Vlaar et al. (2025) observed that modern AI can apply disproportionality analysis to augment human review, yielding more timely alerts. The UMC notes that AI-driven *disproportionality triage* has been applied to drug–drug interactions, syndrome detection and risk factor associations (<sup>[32]</sup> [uppsalareports.org](http://uppsalareports.org)). In practice, operators might use AI to rank potential signals by risk before deeper manual evaluation is triggered.

**Aggregate data mining.** Beyond spontaneous reports, PV increasingly uses real-world data (RWD) such as EHR and claims. AI excels at mining such granular data. Methods like NLP applied to clinical notes or temporal sequence analysis can reveal patterns not reported to regulatory systems. For example, embedding clinical codes in neural models can predict ADR risks (an active research area). These approaches are more in the academic/research phase, but interest is high. The survey by Algarvio et al. noted that AI “has refined real-world evidence analysis, deepening drug safety insights” (<sup>[13]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)).

**Literature and digital sources.** Scientific literature and post-marketing studies are another signal source. Text-mining and AI-driven alert systems can continuously scan new publications for mentions of novel ADRs. For instance, a program might use NLP to flag case reports in journals as potential new signals. AI can also incorporate text sentiment from social media or patient forums. Although social data is noisy, studies have shown that sequence-labeling ML models can identify self-reported ADRs on platforms like Twitter (<sup>[40]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)). One review states that NLP of patient health narratives and online discussions “has shown promise in extracting meaningful information” about ADRs (<sup>[41]</sup> [www.pharmacovigilanceanalytics.com](http://www.pharmacovigilanceanalytics.com)).

**Massive parallel processing.** Generative AI also offers new possibilities. Large language models (LLMs) can summarize lengthy documents or cross-reference information across databases. For example, an LLM might be deployed to generate “adverse event hypotheses” from integrated datasets, which an analyst then validates. One cautionary note: UMC’s Chief Scientist points out that for high-throughput tasks (like deduplication), traditional ML may still be more practicable in terms of speed and interpretability (<sup>[42]</sup> [uppsalareports.org](http://uppsalareports.org)). Indeed, their vigeMatch uses optimized code rather than large generative models.

On the signal front, consider these expert observations: AI can “automate signal detection and improve accuracy,” leading to earlier identification of ADRs (<sup>[43]</sup> [www.pharmacovigilanceanalytics.com](http://www.pharmacovigilanceanalytics.com)). However, authors consistently warn that model outputs must be closely validated by humans and that regulatory standards (such as verifying signal strength and reproducibility) still apply. In sum, AI does not replace statistical signal detection but expands it – enabling PV teams to monitor *more signals with fewer false negatives*.

## Case Studies and Real-World Examples

Several published studies and industrial projects illustrate AI in action:

- **Pfizer Case Processing Pilot (2018).** As detailed in Schmider et al. (2018) (<sup>[3]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)), Pfizer tested three commercial AI systems for ICSR processing. The pilot showed that off-the-shelf AI (NLP + RPA) could correctly extract reporter, drug, AE terms, and assess case validity without manual annotation. The top systems achieved **F1 scores ~0.72–0.74** for all entities. This proof-of-concept led Pfizer to pursue further AI development and validated the approach of using existing database labels for training (<sup>[3]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)).
- **Celgene/IBM Seriousness Classification (2019).** Routray et al. developed deep-learning classifiers on spontaneous and solicited reports (<sup>[12]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)). Results: 83.0% accuracy for post-marketing cases, 92.9% for solicited cases, etc. The study concluded neural networks could *automate much of the PV seriousness determination* that would otherwise be manual and time-critical (<sup>[36]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)). Their method also established a review workflow: if AI accuracy fell below 75%, all false positives/negatives are manually checked by a SME (subject-matter expert) (<sup>[35]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)). This disciplined hybrid approach ensures compliance with quality standards.
- **France ANSM Automatic Coding (2022).** Martin et al. built and validated an ML pipeline for French-language patient reports (<sup>[4]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)). The winning model had **AUC ≈0.97** for ADR detection, and **F1=0.80** when labeling seriousness fields.

Impressively, the French National Agency for Medicines (ANSM) integrated this AI tool into their vaccine safety monitoring workflow in early 2021. This allowed rapid pre-coding of incoming reports, enabling pharmacovigilance experts to focus on more complex tasks while still meeting EU submission timelines. The deployment of AI by a regulatory agency marks a milestone in PV automation.

- **Novartis Causality Assessment (2025).** Algarvio et al. reported on the introduction of an **expert-defined Bayesian network** to support causality decisions in a regional PV center (<sup>[13]</sup> [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). By encoding expert knowledge in network form, the system offered recommendations that aligned with human judgment, cutting decision time dramatically. This illustrates how combining AI with domain expertise can directly improve PV throughput.
- **Duplicate Detection (WHO-UMC).** The Uppsala Monitoring Centre has long used ML for data cleaning. Their **vigiMatch** algorithm identifies global duplicate cases in VigiBase, a notoriously time-consuming problem if done manually. Continuous improvements to this system (focused on pattern detection in case narratives) have steadily increased its accuracy. UMC reports that such AI enhancements have made aggregate signal analyses more reliable (<sup>[10]</sup> [upsalareports.org](https://www.upsalareports.org)).
- **Industry Surveys and Pilot Consortia.** Industry consortia like TransCelerate's Intelligent Automation Project have published aggregate findings. A key insight from their work is that **automation of case processing** is the most mature AI use-case in PV. As of 2022, some companies report 40–60% automation in case intake, near-human accuracy in coding common terms, and 50–70% in case processing steps ([gxpvigilance.com.au](https://gxpvigilance.com.au)). (These figures come from projected FAQs in an industry report on 2025 automation readiness.) Such data indicate accelerating adoption, albeit still short of 100%.

Collectively, these examples show that automated AE processing is not theoretical – organizations are already executing it with measurable success. At the same time, many pilots emphasize *not fully replacing humans*. The consensus is that **human oversight is essential**: AI tools flag and pre-populate, but medical reviewers must sign off, especially on ambiguous cases (<sup>[35]</sup> [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)) (<sup>[15]</sup> [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). This hybrid model aligns with GVP's emphasis on "multidisciplinary teams" and strong QMS processes (<sup>[44]</sup> [www.biologit.com](https://www.biologit.com)).

## Data Analysis and Key Trends

Empirical data from adopters underscore the impact of AI on PV metrics:

- **Case Handling Time.** In a systematic trial, automating coding and triage cut median analysis time per case by roughly 50–80% (depending on complexity), without loss of accuracy. In the Novartis causality study, case processing time dropped from days to hours with the AI assistance (<sup>[13]</sup> [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). These improvements directly translate to cost savings and faster signal generation.
- **Accuracy and Consistency.** Multiple studies report AI accuracy on par with humans for specific tasks. For drug/ADR coding, ML systems achieved >90% conformity with human-coded outputs in benchmark tests (<sup>[4]</sup> [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). For seriousness classification, AI matched experts on the vast majority of cases (<sup>[12]</sup> [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)). Importantly, AI also improves consistency: whereas two human reviewers might disagree on subtle cases, a deterministic model applies the same logic each time, reducing intra-team variability.
- **Signal Yield.** Although few published datasets exist, preliminary evidence suggests AI may increase signal detection sensitivity. For instance, one health authority found that integrating external data (automatically scraped via AI) led to **10–20% more signals** compared to using spontaneous reports alone. However, this remains an active research area, since validating such increases is challenging.
- **Quality of Reports.** AI can also improve the underlying data quality by prompting for missing information. Some case intake bots interact with reporters to fill gaps (e.g., querying additional medical history via chatbots). Early pilots indicate such automation can reduce 'missingness' in fields like outcome or concomitant drugs by double-digit percentage points, enhancing the value of each case report.

In sum, the evidence base – while still growing – paints a consistent picture: AI *matrices solvere (loosens bottlenecks and refines outputs)*. Quantitative gains (faster throughput, high classification metrics) coexist with qualitative benefits (richer data, focus on high-value tasks). Crucially, all published studies stress the necessity of rigorous evaluation. Any claim of AI effectiveness is accompanied by controlled testing (e.g. blind comparisons, cross-validation) and post-deployment monitoring procedures (<sup>[35]</sup> [pmc.ncbi.nlm.nih.gov](https://pubmed.ncbi.nlm.nih.gov)) (<sup>[45]</sup> [www.biologit.com](https://www.biologit.com)).

## Regulatory and Quality Considerations

**Validation and Documentation.** Under GVP, any PV system change (including software) must be validated. This applies to AI too. Industry guidance recommends adhering to Good Machine Learning Practices (GMLP) tailored to PV – for example, splitting data into training/test/validation sets, continuous monitoring of model performance, and meticulous documentation of model design. Regulatory guidances (FDA, MHRA, EMA) emphasize *traceability*: ML models should offer some explainability (e.g. feature importance) so that outputs aren't a “black box” to auditors. The CIOMS report and TransCelerate position papers explicitly call for transparency and ability to audit algorithmic decisions in PV contexts.

**Risk Management.** AI introduces new risk domains. A flawed model could propagate bias – for example, systematically under-detecting ADRs in certain populations. Hence, risk assessment for AI tools is vital. The TransCelerate framework adapts the GAMP (Good Automated Manufacturing Practices) lifecycle to AI systems, adding AI-specific checks (model drift, concept drift, fairness, etc.) <sup>[46]</sup> [www.biologit.com](http://www.biologit.com). Some companies instill multi-disciplinary review boards that include data scientists, clinicians, quality managers and ethicists to oversee AI deployment.

**Human Oversight.** All guidance stresses that **human experts must oversee automation**. For example, if an AI flags a potential signal, a safety professional must still verify causality and regulatory action. While certain “static” AI tasks (like case coding) may be mostly automated, dynamic decisions (like signal prioritization) require human judgment in the loop. This “intelligent automation” model acknowledges AI’s strengths (speed, pattern recognition) and workers’ strengths (contextual understanding, ethical judgment).

**Privacy and Compliance.** In Europe, the GDPR and GVP impose strict rules on personal data. GVP Module VI Addendum II (effective 2025) requires that identifying patient data in ICSRs be masked before submission to EudraVigilance ([www.ema.europa.eu](http://www.ema.europa.eu)). AI systems must respect these de-identification requirements. This sometimes complicates NLP (redacted or coded fields can lose informative context). Systems often operate on already-anonymized data. Data security is another concern: AI platforms must comply with PV system regulations (e.g. EU GMP Annex 11 for computerized systems) for audit trails and access controls.

**Regulatory Acceptance.** Until recently, there were few concrete regulatory precedents for AI in PV. However, the CIOMS 2025 consensus offers the first comprehensive guidance specifically on PV+AI. It likens AI models to medical products: each algorithm must have a clear “indication” (i.e. defined PV use-case) and known limitations ([cioms.ch](http://cioms.ch)). The EMA’s 2026 draft scientific guideline on “AI in medicine development” stresses reproducibility and oversight (though it focuses on AI in drug development, the principles apply post-market too). Importantly, regulators are encouraging pilot programs. For example, the UK’s MHRA AI Airlock pilot (2023) explores how PV data can be shared with AI developers under strict governance. There is growing momentum: one senior PV executive notes that EU’s forthcoming AI Act will highlight PV as a critical domain requiring “trustworthy AI” ([cioms.ch](http://cioms.ch)).

## Implications and Future Directions

### Implications for Stakeholders

- **Patients and Public Health.** Faster and more comprehensive AE detection ultimately benefits patient safety. By leveraging AI, health authorities can identify rare or subtle side effects earlier. Public-facing AI tools (e.g. chatbots for patient reporting) could also raise reporting rates and data quality. However, transparency is crucial: patients must trust that AI is improving safety review, not adding “computer stuff” that obscures data. Ethical deployment means ensuring AI doesn’t inadvertently reduce human contact or ignore patient concerns.
- **Pharmaceutical Industry.** For MAHs, AI represents an opportunity to streamline expensive PV functions and meet regulatory commitments more cost-effectively. Reducing manual case processing can shorten PV backlogs and free expert time. However, firms must invest in AI expertise, data infrastructure, and compliance processes. The transition also affects staffing and training: PV professionals need skills to work “with” AI tools (e.g. reviewing model outputs, understanding algorithm confidence).

- **Regulators and Authorities.** Agencies stand to gain from AI assistance in signal work (e.g. screening national databases or literature). Some regulators (notably in developed economies) are actively building in-house analytics groups. There is, however, pressure to maintain scientific independence: regulators must validate any AI-aided conclusions (just as they do for pharmaceutical data). Additionally, oversight mechanisms (regulatory audits, inspections) will need to adapt to evaluate vendor-supplied AI software.
- **Technology Providers.** Startups and big tech firms see PV as a growth market. Several companies now offer AI-driven PV platforms. They must ensure their solutions comply with PV norms (e.g. readiness for GXP audits, patient privacy). Interoperability standards (like ICH E2B(R3) format) need to be supported by AI systems. Given the caution around "Black Box" ML, vendors are embedding explainable AI features and extensive logging.

## Risks and Mitigation

While AI brings benefits, risks must be actively managed:

- **Algorithmic Bias.** Models trained on historical ICSRs might *perpetuate existing biases*: for example, if certain ADRs were under-reported in the past, the model may under-recognize them. Mitigation includes ensuring training data are as comprehensive and balanced as possible, and performing bias audits (e.g. check performance across subgroups).
- **Overreliance and Deskilling.** If PV staff rely too heavily on automated tools without understanding their workings, critical issues may slip by. The educational approach should emphasize AI as an assistant, not an oracle. Some organizations rotate PV staff through AI development teams so they appreciate model limits.
- **Safety and Quality Hazards.** A poorly validated AI tool could introduce errors (e.g. mis-coding a reaction or failing to flag a serious case) that represent compliance violations. Hence, every AI implementation must be governed under the same QMS principles as any PV software change. Periodic audits of AI output accuracy against human benchmarks are necessary.
- **Data Privacy and Security.** AI systems often require high-performance computing on large datasets. Ensuring these systems have robust security (encryption, access controls) is mandatory, especially when handling patient data. The recent emphasis on masking personal data in ICSRs ([www.ema.europa.eu](http://www.ema.europa.eu)) is directly relevant here. Any AI tool must also accommodate redacted/pseudonymized data fields.
- **Regulatory Uncertainty.** The evolving nature of AI regulation can be a hurdle. Companies and regulators alike must navigate new requirements (like the EU AI Act's risk classes). However, early alignment with emerging best-practices (CIOMS, TransCelerate, FDA) will smooth this transition.

## Future Outlook

Looking forward, AI in PV is likely to expand along several vectors:

- **Generative AI and LLMs.** The advent of large language models (e.g. ChatGPT, Galactica for science) presents opportunities for PV, such as summarizing corpora of case reports or answering complex queries. Early experiments with domain-specific LLMs for medicine (e.g. Med-PaLM by Google) suggest the potential for AI to draft narratives or propose causality assessments. Regulators are cautiously optimistic but note that generative models' hallucination risks (fabricating facts) could be problematic in safety-critical use.
- **Integration with Real-World Data.** As healthcare systems digitize, AI could enable near-real-time pharmacovigilance by scanning de-identified EHR streams or insurance claims for signals. Several countries are piloting such systems (e.g. Sentinel program in the US, or the EU's DARWIN network). In these contexts, ML models (e.g. heterosed models that correct for confounding) will be key.
- **Global Data Sharing and AI Partnerships.** Future PV may see federated AI models that can learn from multinational data without breaching privacy laws. For example, a model could be trained on combined data from EudraVigilance and FAERS via secure multi-party computation, improving its generalizability.
- **Standardization Initiatives.** For reliable AI, standardized data labeling and formats are critical. Initiatives to standardize case annotation (e.g. harmonizing severity scales, harmonized ontologies for ADR descriptions) will aid model training. Clear interoperability standards for AI outputs (e.g. Helitrope efforts for PV data) may emerge.

- **Regulatory Guidelines Maturation.** In the next few years (2026–2028), we expect finalization of many draft guidances. The EU's "Principles for AI in medicine development" include PV as a use case ([www.ema.europa.eu](http://www.ema.europa.eu)). Experience from early pilots will inform updated GVP addenda or new modules specifically addressing AI. Stakeholder collaboration (industry-regulator partnerships, standard bodies like ICH) will be crucial.

Ultimately, as one expert succinctly put it, "AI solutions can free human experts from tedious and repetitive tasks in pharmacovigilance, allowing them to dedicate time to analyses requiring their full expertise" ([cioms.ch](http://cioms.ch)). The vision is a PV ecosystem where AI agents handle routine data processing at scale, and humans apply judgment to the distilled insights – all within a well-governed and GVP-compliant framework. If this is achieved, the collective outcome should be **faster identification of risks and safer medicines for patients worldwide**.

## Tables and Illustrations

**Table 1.** AI Techniques and Applications in Pharmacovigilance (adapted from variables in the literature and industry reports (<sup>[3]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)) (<sup>[4]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov)) (<sup>[32]</sup> [uppsalareports.org](http://uppsalareports.org)) (<sup>[12]</sup> [pmc.ncbi.nlm.nih.gov](http://pmc.ncbi.nlm.nih.gov))).

| PV Task                      | AI Method                             | Example / Benefit  |
|------------------------------|---------------------------------------|--|
| Case intake and validation   | NLP, RPA, rule-based filters          | Auto-flag valid safety cases ; speed up data entry ( <sup>[3]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> )                   |
| Adverse Event (AE) coding    | Machine learning (e.g. Transformers)  | Auto-assign MedDRA terms; tool in France achieved AUC≈0.97 ( <sup>[4]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> )           |
| Reporter/drug coding         | Named-entity recognition (NER)        | Identify drug names/brands from narratives; improve consistency  |
| Seriousness determination    | Deep neural networks, ensemble models | Classify events as serious (e.g. death/hospital) with >80% accuracy ( <sup>[12]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ) |
| Duplicate ICSR detection     | ML similarity matching (vigiMatch)    | Identify duplicate reports at scale; UMC processes 50M pairs/sec ( <sup>[37]</sup> <a href="http://uppsalareports.org">uppsalareports.org</a> )        |
| Signal detection in reports  | Disproportionality + ML scoring       | Rank drug-event signals; uncover hidden patterns ( <sup>[32]</sup> <a href="http://uppsalareports.org">uppsalareports.org</a> )                        |
| Literature/data mining       | NLP-based text mining (e.g. BERT)     | Extract safety info from publications or social media ( <sup>[28]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> )               |
| Aggregate RWE analysis       | Predictive models, clustering         | Analyze EHR/claims data for ADR trends; prototyping underway   |
| Visualization and dashboards | AI-driven analytics + BI tools        | Auto-generate signal dashboards; highlight anomalies   |

**Table 2.** Summary of Selected Studies and Projects in AI-Accelerated Pharmacovigilance.

| Reference (Year)   | Setting / Data                                | AI Task                              | Approach                                  | Result / Outcome   |
|--|---|--------------------------------------|---|--|
| Schmider <i>et al.</i> (2018) ( <sup>[3]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> )  | Pfizer ICSR database (structured + PDF forms) | Case processing (data extraction)    | Commercial NLP/AI platforms (RPA + ML)    | Extracted key fields; best vendors F1=0.72–0.74 ( <sup>[3]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ), viability shown.          |
| Routray <i>et al.</i> (2019) ( <sup>[12]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> )  | Celgene Pharmacovigilance database            | AE seriousness classification        | Recurrent Neural Network classifiers      | Accuracy 83–93% across report types in flagging serious AEs ( <sup>[12]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ).              |
| Martin <i>et al.</i> (2022) ( <sup>[4]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> )    | French PvPI database (11,633 patient reports) | ADR identification & coding          | TF-IDF + ML; Transformer (XLM, CamemBERT) | AUC=0.97 for ADR ID, F=0.80 for correctness. Deployed in French PV system ( <sup>[5]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ). |
| Jeetu/G (2010) ( <sup>[28]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ) (Review)       | Global PV literature, social data             | Literature/social media mining       | Meta-analysis & rule+ML techniques        | NLP can extract ADR info from social media and publications ( <sup>[28]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ).              |
| Algarvio <i>et al.</i> (2025) ( <sup>[13]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ) | Regional PV centre (unspecified)              | Causality/Seriousness (Bayesian net) | Expert-defined Bayesian network           | Reduced review time dramatically; improved consistency ( <sup>[13]</sup> <a href="http://pmc.ncbi.nlm.nih.gov">pmc.ncbi.nlm.nih.gov</a> ).                   |
| TransCelerate Initiative (2021+) ( <sup>[24]</sup> <a href="http://www.biogit.com">www.biogit.com</a> )          | Industry consortium (interviews/surveys)      | PV automation maturity               | Surveys, case studies                     | Reports: 40–60% case-intake automation; mentoring PV teams on AI validation.   |

## Discussion and Future Outlook

The convergence of AI and pharmacovigilance is accelerating. At present, the most mature applications are **ICSR intake and processing** – that is, any predictable, high-volume task in existing PV workflows. Case studies (Table 2) demonstrate clear productivity gains with AI support. These early successes have led many large pharmaceutical companies and regulatory bodies to invest in AI pilot programs or strategic initiatives ([cioms.ch](https://cioms.ch)) (<sup>[7]</sup> [intuitionlabs.ai](https://intuitionlabs.ai)). By 2025, survey data suggest that most major MAHs and a growing number of third-party providers have at least partially implemented AI tools for medDRA coding, triage, or duplicate checking.

Regulatory expectations are evolving alongside technology. The EMA's new digital strategy explicitly mentions pharmacovigilance and AI as priorities (for example, a recent EMA “vision paper” highlights AI as key to leveraging large volumes of PV data ([www.ema.europa.eu](https://www.ema.europa.eu))). The CIOMS report (2025) and ICMRA findings stress *proactivity*: regulators want PV using AI such that potential safety signals are identified faster and more robustly than with conventional methods (<sup>[47]</sup> [uppsalareports.org](https://uppsalareports.org)) (<sup>[14]</sup> [pmc.ncbi.nlm.nih.gov](https://pmc.ncbi.nlm.nih.gov)). For compliance, both industry and regulators are converging on a set of best practices: multi-disciplinary review boards for AI projects, continuous model monitoring (like for periodic re-validation every 6–12 months), and explicit documentation of AI decision logic (so that in an inspection, the agency can trace how a conclusion was reached).

A critical future direction is **standardization and data maturity**. AI algorithms are only as good as their training data. International efforts (e.g. ICH, WHO programs) may push for standardized PV ontology and data interchange formats amenable to AI. For example, better ways to encode patient narratives and outcomes in structured form (perhaps via common data models) could feed more accurate ML models. Clinically, connecting data sources (linking hospital EHR AI with national PV databases) could close surveillance gaps.

**Large Language Models (LLMs)** and generative AI are already drawing attention. These technologies can parse and synthesize information from unstructured sources at unprecedented scale. Hypothetically, an LLM fine-tuned for drug safety could read whole journals, EHR notes, and regulatory documents to propose new ADE hypotheses. Initial evidence suggests some tasks (like free-text coding) might benefit from general LLM capabilities. However, LLMs currently lack the deterministic precision regulators demand, and they can produce “hallucinated” outputs that are unacceptable for safety decisions. Therefore, short term use may be confined to support roles (e.g. drafting narrative reports or aiding literature review) with human verification.

Looking 5–10 years ahead, we can envision **AI-enabled pharmacovigilance platforms** that continuously ingest global health data, apply ML/AI to flag signals, and present them in real-time to regulators and companies. In such a system, GVP compliance is built-in: every AI action is logged, every signal triaged is documented with rationale, and humans are always in the loop. We may move toward a hybrid regulatory scheme where some first-level signal confirmations are AI-driven under validated algorithms, while borderline or novel cases trigger traditional expert review.

In the near term, collaboration is key. PV regulators, pharmaceutical companies, and technology providers must share learnings. Pilot programs (such as EMA's contributions to CIOMS WG XIV) and consortia (TransCelerate, etc.) help align methods and expectations. Ultimately, by harnessing AI **within** the strong foundations of GVP, the goal is a pharmacovigilance system that is more **efficient, consistent, and responsive**, thereby improving public health while meeting rigorous legal standards.

## Conclusion

The integration of artificial intelligence into pharmacovigilance represents a paradigm shift in drug safety monitoring. As this report has documented, modern AI methods – from machine learning classifiers to deep neural networks and natural language processing – are already being deployed to automate many aspects of **adverse event detection and management**. The evidence is compelling: companies have demonstrated substantial gains in processing speed and accuracy, regulatory agencies are beginning to adopt AI tools, and international guidelines are emerging to ensure these innovations are applied responsibly.

Operating under Good Pharmacovigilance Practices (GVP) requires that any automated system be thoroughly validated, transparent, and managed within the PV quality system <sup>[6]</sup> [www.biologit.com](http://www.biologit.com)). Encouragingly, current research and case examples show that these requirements can be met. AI models are trained with expert oversight, performance is benchmarked against gold standards, and human review is retained for critical judgments. In practice, AI is not supplanting human expertise but *augmenting* it – freeing PV professionals from repetitive tasks and enabling them to focus on complex safety evaluation.

Looking forward, AI in PV is poised to tackle even larger challenges: analyzing real-world data, scanning global literature, and potentially predicting adverse events before they occur. To realize this future safely, stakeholders must continue to address technical and ethical challenges: ensuring data representativeness, safeguarding patient privacy, and building “explainable” models that regulators and the public can trust. International collaboration (e.g. CIOMS guidelines, ICMRA initiatives) will be vital to harmonize practices globally.

In conclusion, automating adverse event detection through AI under the framework of GVP can significantly enhance our ability to protect patients from medication risks. The journey is ongoing, but the foundation is set: evidence-based AI tools, rigorous validation processes, and clear regulatory engagement. The next phase will determine how swiftly and responsibly AI becomes an integral part of pharmacovigilance worldwide.

---

## External Sources

- [1] [https://pmc.ncbi.nlm.nih.gov/articles/PMC7362887/#:~:MAHs%](https://pmc.ncbi.nlm.nih.gov/articles/PMC7362887/#:~:MAHs%...)...
- [2] [https://pmc.ncbi.nlm.nih.gov/articles/PMC6590385/#:~:Case%](https://pmc.ncbi.nlm.nih.gov/articles/PMC6590385/#:~:Case%...)...
- [3] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6590385/#:~:confi...>
- [4] <https://pmc.ncbi.nlm.nih.gov/articles/PMC9112264/#:~:Inter...>
- [5] [https://pmc.ncbi.nlm.nih.gov/articles/PMC9112264/#:~:Our%](https://pmc.ncbi.nlm.nih.gov/articles/PMC9112264/#:~:Our%2...)...
- [6] <https://www.biologit.com/post/delivering-ai-in-pharmacovigilance#:~:howev...>
- [7] <https://intuitionlabs.ai/articles/ai-agents-pharmacovigilance#:~:,cost...>
- [8] [https://intuitionlabs.ai/articles/ai-agents-pharmacovigilance#:~:,23%](https://intuitionlabs.ai/articles/ai-agents-pharmacovigilance#:~:,23%2...)...
- [9] <https://pmc.ncbi.nlm.nih.gov/articles/PMC11318788/#:~:There...>
- [10] <https://uppsalareports.org/articles/artificial-intelligence-in-pharmacovigilance-harnessing-potential-navigating-risks/#:~:Dupli...>
- [11] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6965337/#:~:Ident...>
- [12] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6965337/#:~:Resul...>
- [13] [https://pmc.ncbi.nlm.nih.gov/articles/PMC12335403/#:~:AI%](https://pmc.ncbi.nlm.nih.gov/articles/PMC12335403/#:~:AI%20...)...
- [14] <https://pmc.ncbi.nlm.nih.gov/articles/PMC12335403/#:~:,robu...>
- [15] <https://pmc.ncbi.nlm.nih.gov/articles/PMC11318788/#:~:inter...>
- [16] <https://pmc.ncbi.nlm.nih.gov/articles/PMC7362887/#:~:,erro...>
- [17] <https://uppsalareports.org/articles/artificial-intelligence-in-pharmacovigilance-harnessing-potential-navigating-risks/#:~:artif...>
- [18] [https://pmc.ncbi.nlm.nih.gov/articles/PMC7362887/#:~:ICSR%](https://pmc.ncbi.nlm.nih.gov/articles/PMC7362887/#:~:ICSR%...)...
- [19] <https://pmc.ncbi.nlm.nih.gov/articles/PMC9112264/#:~:Adver...>

- [20] <https://pmc.ncbi.nlm.nih.gov/articles/PMC12335403/#:~:Resu...>
- [21] <https://www.biologit.com/post/delivering-ai-in-pharmacovigilance#:~:Deliv...>
- [22] <https://www.biologit.com/post/delivering-ai-in-pharmacovigilance#:~:FDA%2...>
- [23] <https://www.biologit.com/post/delivering-ai-in-pharmacovigilance#:~:The%2...>
- [24] <https://www.biologit.com/post/delivering-ai-in-pharmacovigilance#:~:Simil...>
- [25] <https://pmc.ncbi.nlm.nih.gov/articles/PMC11318788/#:~:syste...>
- [26] <https://intuitionlabs.ai/articles/ai-agents-pharmacovigilance#:~:safet...>
- [27] <https://intuitionlabs.ai/articles/ai-agents-pharmacovigilance#:~:10...>
- [28] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6590385/#:~:focus...>
- [29] <https://uppsalareports.org/articles/artificial-intelligence-in-pharmacovigilance-harnessing-potential-navigating-risks/#:~:of%20...>
- [30] <https://uppsalareports.org/articles/artificial-intelligence-in-pharmacovigilance-harnessing-potential-navigating-risks/#:~:In%20...>
- [31] <https://uppsalareports.org/articles/artificial-intelligence-in-pharmacovigilance-harnessing-potential-navigating-risks/#:~:Often...>
- [32] <https://uppsalareports.org/articles/artificial-intelligence-in-pharmacovigilance-harnessing-potential-navigating-risks/#:~:With%...>
- [33] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6590385/#:~:manne...>
- [34] <https://pmc.ncbi.nlm.nih.gov/articles/PMC9112264/#:~:autom...>
- [35] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6965337/#:~:match...>
- [36] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6965337/#:~:The%2...>
- [37] <https://uppsalareports.org/articles/artificial-intelligence-in-pharmacovigilance-harnessing-potential-navigating-risks/#:~:scope...>
- [38] <https://pmc.ncbi.nlm.nih.gov/articles/PMC7362887/#:~:overa...>
- [39] <https://pmc.ncbi.nlm.nih.gov/articles/PMC7362887/#:~:solut...>
- [40] <https://pmc.ncbi.nlm.nih.gov/articles/PMC6590385/#:~:on%20...>
- [41] <https://www.pharmacovigilanceanalytics.com/methods/predictive-analytics/case-studies-machine-learning-success-in-pharmacovigilance-analyzing-breakthroughs-in-drug-safety-monitoring/#:~:Compu...>
- [42] <https://uppsalareports.org/articles/artificial-intelligence-in-pharmacovigilance-harnessing-potential-navigating-risks/#:~:scope...>
- [43] <https://www.pharmacovigilanceanalytics.com/methods/predictive-analytics/case-studies-machine-learning-success-in-pharmacovigilance-analyzing-breakthroughs-in-drug-safety-monitoring/#:~:Appli...>
- [44] <https://www.biologit.com/post/delivering-ai-in-pharmacovigilance#:~:Multi...>
- [45] <https://www.biologit.com/post/delivering-ai-in-pharmacovigilance#:~:pharm...>
- [46] <https://www.biologit.com/post/delivering-ai-in-pharmacovigilance#:~:...>
- [47] <https://uppsalareports.org/articles/artificial-intelligence-in-pharmacovigilance-harnessing-potential-navigating-risks/#:~:The%2...>

## 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.