AI-powered pharmacovigilance: From reactive reporting to predictive risk

Aurelia Bordes
Global Head
Safety Operations

Brian Kaulback
Global Head
Technology Solutions – Safety & Logistics

Evan Lin
Global Head
AI Labs

Kausik Maiti
Executive Director
Safety Medical Sciences

Sanjay Vyas
President
Safety & Logistics; Country Head of India

Pharmacovigilance (PV) is entering a new era that will transform safety conduct from today's reactive reporting to predictive risk assessment. This transformation will be fueled by unprecedented volumes of real-world health data and operationalized by artificial intelligence (AI) and predictive analytics. In this paper, we discuss current, near-term and future AI applications, with examples of Parexel's emerging AI solutions to advance efficiencies and insights in pharmacovigilance practice.

Introduction

Pharmacovigilance (PV)—arguably the most labor-intensive task in the drug development continuum—is the ideal workspace for AI technologies. The first wave of AI implementation is increasing automation and streamlining workflows in adverse event case processing and reporting. Increasingly sophisticated AI tools also promise to improve patient safety by enabling deeper insights in safety signal detection based on increasing volumes of real-world data. With the integration of predictive analytics, the future of AI-powered systems will be used to predict risks due to adverse drug events—initially for specific populations, and ultimately for individual patients. Parexel pharmacovigilance experts envision a future in which safety data may be used to support the selection of optimal candidates for drug development, the design of clinical trials, and the practice of personalized medicine.

The immediate challenge is to design AI systems that can scale with growing volumes of safety data while improving quality performance of the high-risk, high responsibility work of safety surveillance. The size of this task—and its demands on pharmacovigilance teams to generate, interpret and monitor adverse event reports—is expanding exponentially.

The challenge: A tidal wave of data

Beginning in clinical trials and extending into real-world use, the safety database captures adverse event (AE) case reports for the entire lifecycle of a drug. Traditional AE source data focuses on spontaneous reports from patient and healthcare professionals, from solicited reports including clinical trials, and from published medical literature.

Expanding data sources

During the past decade, data sources have exploded with access to electronic health records (EHRs), claims databases, patient and disease registries, and social media. At the same time, global marketing has expanded product use from the mature markets of North America, Europe, and Asia Pacific to emerging world-wide markets in South America, Africa, India and Asia. Country regulators are becoming more stringent in their reporting requirements and more focused on aligning requirements with advancing safety surveillance approaches.

Growing data sources and safety awareness are creating dramatic increases in adverse event reports. In 2010, FDA's Adverse Event Reporting System (FAERS) received just over 672,640 AE reports; in 2020, the total exceeded 2.2 million; in 2024, the number escalated to more than 20 million.[1]

A curse and a blessing

The growth of AE sources and monitored populations—together with diverse languages and regulatory reporting formats—have catapulted pharmacovigilance into an exhaustive global mission. While the amount of data to be collected and processed poses a Herculean labor, the deluge of information also holds new levels of understanding—about the effects of therapies in real-world use, about the effects in given populations, and ultimately about the probable effects in an individual patient. The power, both to manage and mine growing volumes of data, is supported by automated processes aided by artificial intelligence.

From 2010-2024, the FDA's Adverse Event Reporting System (FAERS) received:

2024:
20+ million
adverse event reports

2020:
2.2+ million
adverse event reports

2010:
672,640
adverse event reports

AI comes of age: Generative AI and natural language processing

For more than 50 years, PV has depended on spontaneous reporting systems (SRS), like FDA's FAERS, WHO's VigiBase, and EMA's EudraVigilance. These systems are mainstays in signal detection, causality assessment and standardized MeDRA coding. While SRS systems are effective in identifying emerging, unexpected safety issues, they have two critical limitations. They rely heavily on spontaneous reporting, a passive approach that results in under-reporting; it is estimated that 90% of AEs go unreported.[2,3] And they are not designed to manage the unstructured data that is expected to become a major source of AE reports.

Evolving AI solutions

Efforts began in the early 2000s to explore the use of AI to reduce labor and increase speed in pivotal PV tasks like signal detection. Today, signal detection is a semi-automated process. A 2022 review identified 66 articles, published between January 2015 and July 2021, containing information on AI use in pharmacovigilance. Articles focused primarily on use of machine learning to identify AEs; process safety reports; extract drug-drug interactions; identify populations at high risk for drug toxicity; and predict side effects.[4] A 2018 pilot study confirmed the feasibility of utilizing AI to streamline case processing, PV's most resource-intensive task.[5]

As drug developers ascend the AI learning curve, pharmacovigilance is an ideal testing ground. PV case processing and reporting is a stringent, labor-intensive process, flowing from AE case intake; through quality-controlled procedures to assess causality, seriousness and expectedness; and culminating in case report writing and regulatory submission. These well-defined PV processes offer opportunities to automate tasks, increase speed and quality and reduce costs. Current PV applications focus on two AI technologies—generative AI (genAI) and natural language processing (NLP).

GenAI and NLP

GenAI is a type of machine learning that generates new content based on algorithms trained to recognize patterns, build classifications and make predictions. Systems can be trained to recognize drug names, adverse events, or potential patient designations. NLP systems analyze natural language in text and speech data. These models are used to provide translations, summaries, name and entity recognition, and text prediction (like predicting the next word in a sentence), and sentiment analysis (determining whether a text's emotional tone is positive, negative or neutral).

On the horizon: Large Language Models (LLMs)

LLMs are next-generation natural language processing tools that can be trained on large volumes of data to focus on understanding and generating human language. These advanced tools perform more complex data analysis to generate text, provide translations, and answer questions.

Great expectations: AI's potential in PV

In addition to freeing PV teams from time-consuming tasks, increasingly sophisticated AI systems can identify more subtle safety signals and causal relationships in complex data than human reviewers can detect.

Wider, deeper literature review

While conventional automation is limited to searching the medical literature for specific terms, AI systems can learn to recognize relationships between drugs and documented AEs and discern their relevance. Emerging AI applications could assess the probability of an AE in any given citation, prioritize the most critical reports for reviewers, and automatically order articles and translations. By identifying and directing reviewers to the most relevant points in a sea of information, AI is beginning to deliver compelling efficiencies; a 2025 review cites time reductions of 50%.[6]

Better signal detection

AI systems will be able to learn to distinguish among safety signals by comparing signals to underlying reported cases. For example, AI could determine if an adverse event related to a drug combination is an unknown interaction, and if it is, to prioritize it. AI could also be trained to determine whether a signal is an expected interaction for a given population. By cutting through signal noise and reducing false positives, AI will allow PV teams to evaluate a greater number of probable signals rather than limiting their attention to signals indicating the greatest healthcare impact.

Parexel initiatives: Bridge to the AI future

It is tempting to imagine a push-button AI system that completely automates all PV activities. Pharmacovigilance bears too much responsibility for patient safety to surrender its functions to total automation. Human expertise and oversight must be built into operations as sponsors / MAHs transition to AI's greater efficiencies and deeper insights.

In pharmacovigilance, Parexel is developing novel AI applications in three targeted areas: 1) *search and retrieval*, to extract and leverage targeted terms in literature searches and AE source data; 2) *workflow automation*, to improve efficiencies in the time-intensive, repetitive processes in safety case intake and processing; 3) *content authoring*, to accelerate high-quality, comprehensive text creation in medial writing, narratives and reporting documentation.

Together with our partners, Parexel has adopted an AI technology platform that integrates the entire case intake process—from incoming source data to processing in safety databases, through AE assessment, to report writing. These tools are demonstrating impressive benefits.[7]

One size does not fit all

While the potential benefits are enormous, AI is not a quick fix for safety management. AI's power is based on accumulated knowledge specific to a task and the user environment.

AI-first customization

In Parexel's experience, a single AI model cannot be relied on to provide successful results. Overlaying AI in legacy automated systems is rarely successful. To make models work in real-world PV applications, AI systems must be fine-tuned with specific engineering, workflow customization and extensive data cleaning. Parexel is pursuing AI-first approaches and fit-for-purpose designs that integrate human expertise and user training to support ideal outcomes in the transition to AI-enabled pharmacovigilance.[8]

Human-in-the-loop AI

As adoption moves forward, AI systems must be integrated with human expertise and experience to oversee safety activities. Parexel's human-in-the loop approach combines AI with human expertise to leverage AI in situations where sensitivity demands human accountability or where complete automation is not yet feasible. Parexel has been applying AI in user-centered workflows augmented, but not completely controlled, by AI systems. Decision loops include human experts who review AI outputs using platforms designed fit-for-purpose and evaluated using a combination of historical examples, qualitative experience and standard quantitative ML metrics.

AI-enabled pharmacovigilance circa 2030

In the near future, Parexel's pharmacovigilance experts see rapid progress in AI applications for content generation and workflow automation. Advances will be driven by increasing interoperability of real-world data. LLMs are not yet usable for real-world-data processing. But emerging AI systems will rely on LLMs to enabling greater integration and extraction of information from diverse, relatively unstructured sources like EHRs.

Parexel's literature review tool

Today's AI-enabled literature review has reduced time-to-completion to a single day, but significant manual labor is still required. The future for literature review is AI-automated integration with inbound search databases and outbound directly to the safety database. Integration of LLMs will allow these systems to handle all types of case intakes—to identify, retrieve and categorize safety information across any source document (structure or unstructured), then categorize it. AI systems will be able to automate the entire intake, processing and submission process with human oversight.

Parexel's expectedness tool

In case processing, expectedness determination is a complex, manual task aimed at finding whether a reported AE is a known expected event for the product being monitored or new unexpected information. Determination is based on product labeling, which vary by product name and country, and serve as reference documents. In a given case, there might be 15 events reported for patients taking dozens of different drugs, all reported in multiple countries. Parexel's AI tool manages all relevant information and makes it available on a single display window to support the expert's decision-making. Expectedness assessment time is reduced by 30%; the expected related errors were reduced by 60%. Expectedness data are instrumental in defining continued risk/benefit ratio of monitored drugs, so AI-driven improvement will support enhanced patient safety.

AI-managed workflows

AI will be embedded in end-to-end, semi-automated workflows, freeing PV teams to concentrate on scientific analysis and oversight. Hybrid human-AI solutions will advance tasks that require content generation. Parexel is currently testing genAI models to generate case report narratives, and safety periodic reports authoring which will be the next big field of opportunity.

Narrative creation

Case narratives summarize the sequence of events for a given AE, based on all the information received from that event reports. PV experts prepare narratives for regulatory submissions using templates available within or outside of the safety database. Narratives are written, case by case, and even with pre-populated templates, require significant human intervention to ensure the narrative correctly represents the sequence of events and makes medical sense and therefore continue to be labor-intensive. Parexel is developing a tool utilizing LLMs to auto-generate a well-constructed and medically accurate case narrative, based directly on the source documents received, rather than relying on data previously entered in structure fields which can include mistakes. Pilots show promising results in content authoring based on a template. Overall, they show improved efficiencies of about 29.3% considering all cases and 48% in more complex narratives.

AI adoption: Headwinds and tailwinds

The primacy of patient safety is a driving force in AI adoption. AI capabilities with potential to reduce drug-related risks and to advance knowledge about drug impacts in real-world use are compelling opportunities for both sponsors and regulators. As always, the same concerns for patient safety pose hurdles for AI applications.

Regulatory acceptance

A key task for regulators is the validation of AI-systems to ensure that safety and quality standards are maintained. Currently, there is no clear validation method; side-by-side evaluation will proceed for human-AI hybrid systems, and AI-managed processes and regulatory reports will be phased in, stepwise. The simplest automated systems will find acceptance first, paving the way for AI management of more complex tasks.

Regulatory working groups are making some progress—recent guidance has been issued by FDA and EMA. Both have expressed the need for validated, transparent systems that can be adequately audited and enhance patient safety without introducing new risks, including biases or errors in signal detection. The EMA's reflection paper on the use of AI, recommends involving regulators in consultations and performing regulatory impact assessments before implementing AI tools in pharmacovigilance.

Similarly, the FDA highlights the importance of compliance with existing regulations, such as GDPR and HIPAA, and the need for updated post-market surveillance guidelines to accommodate AI-enhanced monitoring. FDA discussion papers highlight the potential of AI/ML to transform drug development and manufacturing but also address challenges such as ethical and security considerations, improper data sharing, cybersecurity risks, and algorithmic discrimination. As the field progresses, the ongoing dialogue between regulators, industry stakeholders, and AI experts will be critical in shaping the future landscape of AI-driven drug safety monitoring and development.

Sponsor acceptance

At Parexel, we see strong interest in AI applications to accelerate timelines and optimize performance. We also see reluctance inspired by concerns about regulatory acceptance, and data privacy. Sponsors recognize the ethical & data privacy issues involved in sharing the huge amounts of data and the level of detail required by AI systems, especially LLMs. They want assurance of rigorous ethical standards and strong firewalls. And they are reporting regulatory scrutiny; inspections now include the query, "Are you using LLM in any of your processes?"

To ensure the highest standards, Parexel governs its AI adoption according to six principles:[9]

- Thoughtful design and deployment
- Accountability and senior-level governance
- Human oversight and control
- Transparent AI
- Regulatory conformance and legal compliance
- Security and privacy

Quo vadis: Envisioning the future of safety practice

On a more distant horizon, two trends point to advancing capabilities and applications of pharmacovigilance: continuous monitoring of real-world patient data; and growing confidence in predictive analytics. Parexel experts envision transformative uses of AI-powered pharmacovigilance data—from the earliest point in drug evaluation to decision-making in the clinic—that will dramatically improve patient safety.

Increasing a focus on the patient

Patient experience is the beginning and end of therapeutic evaluation, and real-world data captures the patient voice in unprecedented ways. AI-powered safety monitoring and reporting will bring patient experience to bear on understanding of unmet medical needs, selection of candidates for development, and patient-focused clinical trial designs.

Informing molecule selection

AI-enabled safety data will come to play a role in molecule selection as genomics and quantum computing advance. Some sponsors are already investigating the potential for evaluations of safety and efficacy in computer-based analyses and simulations. Such *in silico* studies would try the molecule's effects on the human genome in digital format to understand the probability of toxicity and therapeutic effect before human testing. AI-managed safety profiles of similar molecules would inform digital evaluation.

Advancing personalized medicine

The integration of genAI and predictive analytics will create AI systems that confidently predict safety trends and support prevention of harm to patients. Clinicians will be able to predict the likelihood that a drug will cause an adverse event in an individual patient. Drug safety profiles, developed from analyses of huge volumes of real-world, AI-managed data, will be used together with biomarkers in a patient's genetic profile to determine the optimal drug and dosage. These data also may be applied to the development of "digital twins"—simulated patients used as controls in clinical trials.

Strategic PV in the value proposition

The AI-powered pharmacovigilance of the future promises to transform adverse event reports into valuable intelligence for use in developing safer, more patient-focused therapies. Sponsors will be able to use safety data strategically—to align product safety and patient profiles with medical needs and patient experience.

Safety practice will foster deeper understanding of the risks and benefits of marketed products, and to predict the value of innovative therapies in real-world medical practice.

As AI and data integration advance, we envision a future where daily, proactive health monitoring and management becomes the norm, shifting away from current systems singularly focused on pharmacovigilance case processing.

Wearable devices, integrated healthcare provider systems, and comprehensive biopharmaceutical databases will work in concert to create a holistic approach to patient safety. This interconnected ecosystem will enable personalized disease prevention strategies and treatment plans that can proactively identify risk factors by patient populations and therapy types, among other variables to predict outcomes for new therapies. By leveraging these technologies, we may be able to significantly reduce the occurrence of adverse reactions, shifting the focus from reactive reporting to predictive risk assessment. Imagine a world where the current pharmacovigilance practices are no longer needed, and instead the scope and methodologies of pharmacovigilance are redefined, placing greater emphasis on data-driven, AI-powered predictive analytics to – manage upfront risk and ensure patient safety in an increasingly personalized healthcare landscape.

References

1. FDA Adverse Event Reporting System (FEARS): Public dashboard. At: https://www.fda.gov/drugs/fdas-adverse-event-reporting-system-faers/fda-adverse-event-reporting-system-faers-public-dashboard
2. L Hazell, SA Shakir, 2006. Under-reporting of adverse drug reactions: a systematic review. Drug Saf. 29(5):385-396. At: https://pubmed.ncbi.nlm.nih.gov/16689555/
3. P Garcia-Abeijon, C Costa, m Taracido, et al, 2023. Factors associated with underreporting of adverse drug reactions by health care professionals: a systematic review update. Drug Saf. 46(7):625-636. At: https://pubmed.ncbi.nlm.nih.gov/37277678/
4. M Salas, J Petracek, P Yalamanchili, et al, 2022. The use of artificial intelligence in pharmacovigilance: a systematic review of the literature. Pharmaceut Med, 36(5):295-306. At: https://pubmed.ncbi.nlm.nih.gov/35904529/
5. J Schmider, K Kumar, C LaForest, et al, 2018. Innovation in pharmacovigilance: use of artificial intelligence in adverse event case processing. Clin Pharmcol Ther, 105(4):954-961. At: https://pmc.ncbi.nlm.nih.gov/articles/PMC6590385/
6. S Abogunrin, JM Muir, C Zerbini and G Sarri, 2005. How much can we save by applying artificial intelligence in evidence synthesis? Results from a pragmatic review to quantify workload efficiency and cost savings. Front Pharmacol,16:14542-14545. At: https://www.researchgate.net/publication/388569240_How_much_can_we_save
7. Parexel. Putting AI to work in your safety program. At: www.parexel.com/insights/blog
8. A Ghodsi, E Lin, S Vyas. Human-in-the-loop AI: a proven approach to streamlining PV case processing and review. Parexel. At www.parexel.com/insights/article/
9. Parexel. Parexel principles for artificial intelligence. At: www.parexel.com/application/files/resources/assets/Parexel_Principles%20for%20artificial%20intelligence.pdf

With Heart™

We're always available for a conversation

Our solutions are built to help you address today's evolving clinical development landscape and innovate for the future. Highly skilled talent, flexible resources, and guidance from our industry-leading outsourcing experts will provide you with the resilient clinical development strategies and customized models needed for success.

Connect with us us to learn more.

Parexel International Corporation
2520 Meridian Pkwy, Durham, NC 27713, USA
+1 919 544-3170
info@parexel.com

Offices across Europe, Asia, and the Americas
www.parexel.com

parexel®