The red pill alert

Technology is set to play a larger role in monitoring the safety of marketed drugs.

GAURI KAMATH

To effectively monitor drug safety, pharmacovigilance experts are turning to artificial intelligence.

Several years ago, Dhananjay Bakhle received reports of angioedema, an unexpected adverse event, in patients on a medicine marketed by his multinational employer. As the head of drug safety, Bakhle was understandably concerned. Angioedema is a rapid swelling of skin layers, which could be a precursor to anaphylaxis. "It is a very serious side-effect." Global regulators expect to be informed of such serious adverse events (SAEs) within a deadline. (SAEs are suspected, not confirmed, side-effects.) Would he have to flag an entirely new 'safety signal' for a drug marketed worldwide? (A safety signal is new information, on an unexpected adverse event or side-effect potentially caused by a drug, which warrants deeper inquiry.)

No, it turned out. In an attempt to standardise reporting, adverse events are assigned codes as per a specific disease classification or dictionary system. In this case, what was originally reported as oedema (fluid retention) had been wrongly coded, he recalls. Bakhle cites this manual error as one of the many challenges involved in 'pharmacovigilance', or tracking the safety of drugs and vaccines, that technology can help mitigate.

Pharmacovigilance is an expensive, workforce-intensive activity that costs the pharma industry billions of dollars each year.

Pharmacovigilance is a "duty of care" responsibility assigned to companies, says Bakhle. The activity is "so regulated that there are routine inspections by authorities" to check that a company's systems are compliant.

It is also an expensive, workforce-intensive activity that costs the industry billions of dollars each year. Over the past decade or more, the operations have been mostly outsourced to relatively low-cost centres such as India. Thousands of skilled workers, including doctors and pharmacists, are now directly or indirectly employed by the industry to collate, classify, record, and code reported side-effects and to establish 'causality' – how likely it is that the adverse event is caused by the drug.

But increasingly, companies grapple with "an exponential growth of data (on adverse events) from multiple sources," says Selvakumaran Mannappan, Chief Operating Officer, Birlasoft, a Pune-based technology services company which counts pharma companies among its clients. This has led companies to look to artificial intelligence (AI) for solutions.

When it comes to pharmacovigilance, the drug industry has two perennial questions. "How are we going to effectively monitor drug safety and how do we meet evolving regulatory requirements?" says Selvakumaran. And the answer does not seem to lie in throwing more people at the problem. For one, qualified staff have always been in short supply, according to a 2023 White Paper on generative AI and pharmacovigilance by the U.S. data insights company IQVIA, which provides a range of services to the drug industry (*bit.ly/Iqvia-GenAI*). And all companies prioritise "productivity and efficiency", explains Pramod Dhembare, Founder and Managing Director, Fidelity Health Services, a pharmacovigilance services provider.

To make up for staff shortages, traditional software coding has been used to automate operations for two decades, says the IQVIA paper. But it was beginning to seem as though "there was nothing more that could be done with traditional code", it observes. "The industry will have far more data and will need AI to derive insights from this data," it adds.

The data problem

Clinical trials to test a new medicine for safety and efficacy are mostly performed on under 3,000 participants, explains Bakhle, now a veteran of the Indian pharmaceutical industry. A side-effect whose incidence is one in 10,000 is unlikely to get captured in such trials and may come to light post-marketing. This incidence is not considered trivial when lakhs of people will take the medicine, he notes.

Reports of adverse events could come from clinical trials, doctors and hospitals, global adverse event databases, published literature, and the patients themselves. More recently, regulators also expect companies to parse social media chatter. Electronic health records, too, are being looked into and so is information through mobile and wearable devices. The result is what the IQVIA paper calls "overwhelming amounts of data". This is not all structured. For instance, drug names entered in a public database such as the U.S. Food & Drug Administration's Adverse Event Reporting System are not always uniform, and there is duplication. But such systems get lakhs of reports each year.

In parallel, the drug landscape has grown increasingly complex with relatively new categories such as biosimilars, off-patent versions of biologics that involve complex manufacturing and storage, as also cell and gene therapies. These need to be monitored closely in real-world setting. Drug safety regulations are dynamic and vary from one region to another. Companies are also expected to detect and flag new safety signals, says Dhembare. "We have to recognise signals early."

Low-hanging fruit

A good part of pharmacovigilance is about performing repetitive tasks. This includes having to assign priorities based on the seriousness of a reported adverse event, avoid duplication, enter all the relevant medical information per case and review the quality of data to spot errors. While the data may have moved from paper to software platforms, the processes are still workforce-intensive.

The safety imperative

133 drugs were withdrawn for safety reasons between 1990 and 2010

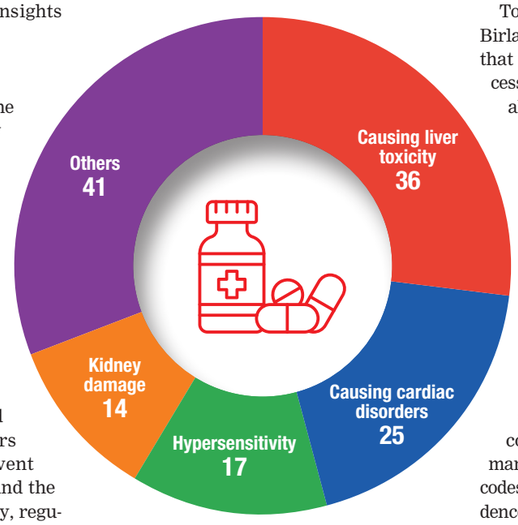

Source: Current Drug Safety, Volume 15, Issue 1, 2020; bit.ly/drugsafety-withdrawal

This could become a bottleneck. At just 7-10 cases per person per day, productivity is the lowest in data entry, says Dhembare. For comparison, a medical reviewer – typically a doctor who assesses causality and writes a medical narrative – can process up to 35 cases a day. A delay in a case entering the software could snowball into a compliance issue, given that serious side-effects have to be reported within a mandated timeline.

The automation of tasks like case intake, data extraction and coding of adverse events can significantly improve productivity and reduce the risk of errors, says Selvakumaran. Birlasoft has developed an AI-powered case intake and triage system for its clients. (Triage involves an initial classification and prioritisation of a case based on seriousness of the reported adverse event.) This could potentially reduce case-processing times by 35%, the company claims. Dhembare of Fidelity has invested in Nuron, a U.S.-based manufacturing automation company. Nuron's data scientists and engineers worked with Fidelity's pharmacovigilance experts to create a soon-to-be-launched process automation tool "with a bit of machine learning thrown in", for case intake and processing.

Pilot studies suggest a 50% cut in processing time with a notable drop in manual errors, Dhembare claims. The halving of the time taken to process a report – from a week to three or four days – is significant as the mandated timeline for reporting a serious adverse event is typically 15 days, he adds. This also saves workforce costs.

To curb manual errors in assigning codes, Birlasoft has developed a generative AI tool that uses advanced natural language processing (NLP) and machine learning (ML) algorithms to process clinical records, research papers, healthcare databases and other data sources. It then generates precise 'ICD' code-mapping by analysing the context and semantics of the provided indications, the company says. (ICD is the International Classification of Diseases system created by the World Health Organization. It is widely used to code adverse events). This aids accuracy, helps categorise patients and facilitates analysis, claims Selvakumaran. The company reports a 75% reduction in manual effort in identifying the necessary codes for data collection from real-world evidence. The tool enhances mapping accuracy and ensures uniformity across datasets, he claims. It also adapts to updates and changes in the ICD standards, he adds.

While faster, less error-prone and more efficient case-processing indirectly helps the cause of safety, its main objective is to improve productivity, free up resources for more strategic tasks, and meet the compliance bar. "This is the low-hanging fruit for AI in pharmacovigilance," says Bakhle.

However, there are other ways in which the tech could directly impact drug safety.

Sending the right signals

There is much that AI can do in the area of "interpretative" processes, such as medical review and narrative writing that form a basis for detecting safety signals, observes Bakhle. This type of writing, done by a

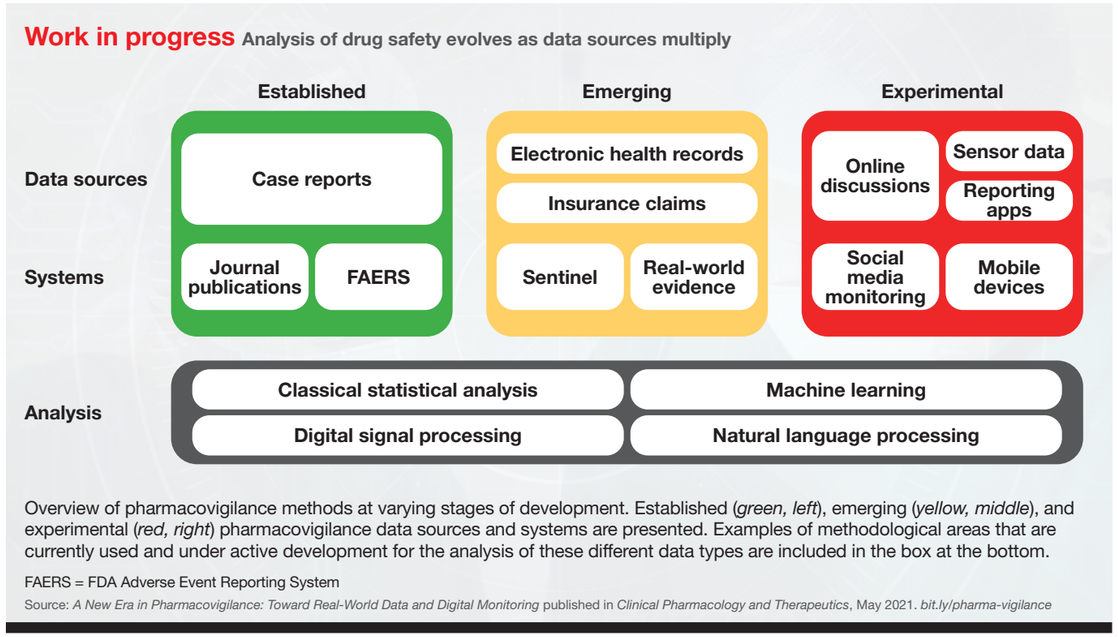

Overview of pharmacovigilance methods at varying stages of development. Established (*green, left*), emerging (*yellow, middle*), and experimental (*red, right*) pharmacovigilance data sources and systems are presented. Examples of methodological areas that are currently used and under active development for the analysis of these different data types are included in the box at the bottom.

FAERS = FDA Adverse Event Reporting System

Source: *A New Era in Pharmacovigilance: Toward Real-World Data and Digital Monitoring* published in *Clinical Pharmacology and Therapeutics*, May 2021. *bit.ly/pharma-vigilance*

medical professional employed in pharmacovigilance operations, involves a narration of all the relevant medical information on a case and concludes with a causality assessment (such as whether the relationship of the adverse event to the drug is 'remote' or 'definite').

Narratives are written in English, and "everybody's English is different", he says citing different ways to describe, say, dizziness. A ChatGPT-like interface could act as a filter – checking each narrative and standardising it in the background, he adds. This would be valuable for a global company that operates in a hundred different countries and has a drug being taken by millions. "Processing is much easier and finding (safety) signals becomes a breeze." NLP techniques can potentially also be used to auto-generate narratives from unstructured data such as medical literature, electronic health records and social media, adds Dhembare.

The IQVIA paper sees AI providing early signal detection, noting that there are already AI tools that can review X-rays faster and more effectively than humans. In 2021, South Korean researchers reported that ML algorithms detected more new signals than conventional methods of analysis in a study. They analysed spontaneously reported adverse events from the Korean government database from two cancer drugs over a nine-year period (*bit.ly/korea-drugsafety*).

Then, in May 2024, researchers from the drug company AbbVie also published an original research article (*bit.ly/AbbVie-drug-safety*) using ML-based approaches to detect potential safety signals for two of its marketed products. It concluded that the model "demonstrated acceptable accuracy for safety signal detection and potential for earlier detection when compared to humans".

Birlasoft has deployed an AI-driven signal-management solution for a pharmaceutical client that was "struggling to harness all the levels of data that they have", says Selvakumaran. By integrating and analysing information from multiple sources, companies can "enhance signal-detection capabilities and perform more comprehensive risk assessments", he adds.

So, could AI help take an unsafe drug off the market? IQVIA's paper identifies some caveats to large-scale adoption. Some of them have been raised for the use of AI in general, such as the quality of data used to train models, the avoidance of biases and the high cost of adoption. The paper also cautions against exposing patient data when training and using these models. It stresses the need to "keep a human in the loop" to maintain control.

As things stand, relying too much on a technology that is evolving could result in short-sighted decisions, cautions Selvakumaran. Any decision related to a drug sold globally would have to be run through multiple models in multiple markets and will be subject to local regulations, he explains. "Speed cannot be a substitute for accuracy," he says. "The data-related analysis that we are trying to do should be an additional input to an expert rather than being an expert." The AbbVie paper adds: "Expert judgment, flexibility, and critical thinking are essential human skills required for the final, accurate assessment of adverse event cases."

Dhembare of Fidelity sees continued advancement of automation and AI over the next few years. Progress may be encumbered by the need for validation and compliance given the highly regulated nature of the sector, he says. But he anticipates "significant improvements" in adverse event detection, signal prioritisation and risk assessment. "This will lead to enhanced patient safety." •

A side-effect whose incidence is one in 10,000 is unlikely to get captured in clinical trials on a few thousand participants.