


Technology's Role in Addressing Potential Drug Safety Risks

MARIE FLANAGAN, Offering Management Director, IQVIA Vigilance Detect

ALISON SLOANE, General Manager, IQVIA Vigilance Detect



Evolving safety landscape

Shortly before COVID-19 started spreading around the world in early 2020, representatives from the European Medicines Agency published “Pharmacovigilance 2030,” an invited commentary predicting that regulators will engage patients and healthcare professionals much more intensively “to maximize the positive impact of pharmacovigilance on the safe and effective use of medicine.”¹

The authors further predicted that smarter pharmacovigilance collection and reporting through the digitization of healthcare provides an opportunity to access more (and better data), and that pharmacovigilance systems should be designed based on an abundance of data rather than a scarcity of it.

Accelerated, no doubt, by the pandemic, these predictions have come to pass — forever changing the way we monitor and address drug safety risks. We are now seeing the continued democratization of pharmacovigilance through technology, with social media and other tools changing our understanding of drug safety monitoring by providing qualitative insights directly from patients. This change is championed by regulators, with initiatives such as v-safe, the Yellow Card System, and other smartphone-based tools enabling direct patient-to-regulator safety communication, making adverse event reporting accessible and further addressing calls for making real-world evidence a key factor in regulatory decision-making.

We’re now starting to see the contours of a pharmacovigilance environment in which technologies will play a critical role in helping to identify drug safety risk efficiently and accurately. With data available from myriad disparate sources, marketing authorization holders (MAHs) must be prepared to capture, analyse and report safety signals from an increasingly connected world — a costly and resource-heavy process that lends itself well to advanced technologies such as artificial intelligence (AI) and natural language processing (NLP)



to help cut through the noise and ensure patient safety. While not all safety data sources are amenable to automation (and human review will always remain a critical component), these tools have been proven to be effective in drug safety monitoring when applied to the right scenario properly.

This paper discusses some of our key learnings in applying an AI-enabled safety risk detection system to many different scenarios across multiple therapeutic areas and offers recommendations for developing a robust drug safety monitoring program to better comply with regulatory requirements.

Using automation to identify potential risk

Monitoring data for potential safety risk is not something MAHs can avoid, it's a regulatory requirement. Global regulators including the EMA and FDA, require companies to screen company-owned data extending as far the internet and digital media assets under their management or responsibility for any potential report of a suspected adverse event. And when an event is found in any of the data under their remit, they have a very tight timeline to report and process that information. That type of data collection, monitoring, and analytics are exceedingly time-consuming and far from fool proof. Drug safety data now flows from just about everywhere and companies must slice and dice it quickly to comply with regulatory agency timelines.

When applied to large data sets, automation promises to alleviate some of this burden for three main types of purposes: remediation, retrospective review, and prospective review. For example, if a company finds they have missed an adverse event, automation can help perform remediation over an extensive data set, so they can find the risk and report it appropriately. There may also be companies who use an AI-enabled safety risk detection system retrospectively to look over all their data to identify a trend or something of interest that may have been overlooked. Lastly, some companies use this technology to monitor their data proactively or prospectively, helping the product and the company in keeping with good compliance and reducing the resources required to do it on an ongoing basis.

The ability of an AI-enabled safety risk detection system such as IQVIA's Vigilance Detect to identify adverse events is not limited to certain forms or datasets. It does look predominantly at social media and digital information (because that's where the volume resides), but the system also can find potential risks in safety databases, customer relationship management systems, quality management systems, and in medical or scientific literature. Safety information also can be derived from virtual agents,

chatbots, or audio files from commercial or medical call centres, and a proper safety risk detection system should be able to identify risks or safety signals from any type of data from any one of these four primary data sources.

Four primary data sources for post-marketing safety surveillance are social and digital media, virtual AI agents, patients support programs, and customer relationship management (CRM) systems. Collectively, these data sources represent a large component of the whole ecosystem of available data from which potential safety signals can be derived.



Social media or digital data includes all company-owned handles and publicly accessible social data that is gathered from company-sponsored apps and wearables.



Virtual AI Agents source data includes chatbot and virtual AI data taken from person/bot interactions to assess for adverse events.



Patient Support Program data exists in multiple formats, both structured/unstructured and multiple languages.



CRM data looks at data residing in customer relationship management systems retrospectively, scanning open/free text fields for potential adverse events.

At IQVIA, we were curious to determine what percentage of that data was successfully managed by technology versus human intervention. In other words, what tipped the balance in terms of how much human intervention was needed depending on the data source?

Effectiveness of technology in identify safety risks

Over the past decade, IQVIA has worked predominantly with Top 10 pharma companies, leveraging the criteria outlined with great success. Annually, we process more than four million records with the vast majority of those coming from social media channels, followed by CRM and virtual agents.

In reviewing our projects, we found that 78% of all data from **virtual agents** and **CRM systems** were successfully managed by technology or automation technology.^{2,3} This means that no human had to intervene in these cases, because the technology was able to assess and route what it believed to be an adverse event confidently. This is a clear-cut use case for AI and NLP as both virtual agents and CRM systems are constructed in a very specific way to ask questions that yield specific, structured answers. So the data that comes in is structured, focused, and clean. As is the case with all data sources, however, human review is still required because everything the machine finds still needs human validation to confirm the reality of the safety event.

In one example of trying to find unreported events in a chatbot, we were able to process more than 292,000 messages which contained 445 adverse events — a small but significant amount.

Patient support programs, by contrast, achieved 90% efficiencies when using a combination of automation techniques, leaving only 10% of digital records for human review.⁴ We found that AI (NLP) had an effectiveness of 37% on its own and achieved an additional 53% when records were automated using a strategic combination of OCR and basic process automation. The difference between analysing patient support programs and other data sources is the type of automation technology that's used. For virtual agents and CRM, it was just AI NLP. But when we applied that to patient support programs, it only had a 37% effectiveness versus 78% in the others. So to get up to 90%, the team fortified the AI, as well as the rules-based algorithms and OCR because the data that's

coming in is so unpredictable. The unstructured nature of that data requires a combination of multiple automation techniques to yield decent results. It takes a lot of work to get the right combination, but it's a worthwhile process if you can remove 90% of the human effort to review 45,000 records per month.

For one customer, our technology supported the retrospective review of 1.5 million documents and found 5,900 previously undetected and unreported safety events rapidly within a 30-day deadline.

We found that AI (NLP) had an effectiveness of 37% on its own and achieved an additional 53% when records were automated using a strategic combination of OCR and basic process automation.

AI and NLP can only do so much when it comes to **social media**, successfully managing around 67% of the data through AI and de-duplication techniques to achieve its efficiencies.⁵ Even though these numbers show it to be relatively successful concerning automation technology, it's the least successful of all data sources. The reason it's not as effective as the other types of data is because of the nature of social media and the nature of the posts themselves. We must have a way to look at things like social media text, emojis, slang, and colloquialisms. NLP can take care of that, but human intervention is often required as social media posts often contain embedded links, URLs, and audio files that require a human for proper validation. So although it might be the least successful area to apply AI in isolation, a robust safety risk identification system trained for social media is still worthy of deployment because it gets the majority of the noise out of the workflow.

In another example, we worked with a Top 10 pharma company to review more than 7.7 million social posts across 300 data sources and 91 languages in 38 countries. By taking into account things like the number of different languages, regional dialects, and country-specific regulations, the study identified more than 100,000 events of relevance to the safety reporting process.

Key components to successful automation


One of the key learnings in our research was that automation and AI only drive certain source types. They drive the success only so far, but it's not a silver bullet for everything. The absolutely critical component is finding the right blend of automation by pairing basic process automation with people and intuitive dashboard-state analytics, then combining that with AI to enable solid results.

IQVIA's Vigilance Detect (formerly AETracker) has been in production for the last 12 years and has been constantly building a bank of ontologies that support the NLP. Exposed to more data over time, the model has been

re-trained by building a bank of specific terms and patterns that make NLP really successful. Our ever-growing bank of ontologies is one of the keys that drives that success. The ontologies are just safety-specific terms and patterns that the system looks for, but they also identify patterns or word proximity within sentences while also looking for colloquialisms, slang, emojis, misspellings, and the like. The ontologies are thorough and specific, enabling the system to identify potential safety events with precision.

Vigilance Detect is designed to leverage all available automation techniques, from NLP to machine learning to human-assisted machine learning. It is most successful when AI technology is paired with basic automation (such as rules-based algorithms and optical character recognition), rather than simply plugging in AI and seeing what happens. Equally critical is having an end-to-end workflow and audit trail that will route information in a very centralized and standardized manner. Properly fashioned, this can bring a huge amount of efficiency to a workflow. Similarly, basic intuitive dashboards point the reviewer to the exact place they need to look, which

FIGURE 1: IQVIA'S VIGILANCE DETECT





Technology is remarkably effective at keeping noise at a minimum to allow humans to focus their attention where it adds the most value.

minimizes the amount of time (and money) spent — and wasted — on the process. Perhaps most critically, human review will always be necessary to validate the event and make sure it meets reporting criteria.

It's important to note that not everything is amenable to automation. Not all data performs exceptionally well with automation, and solutions are most effective when technology is used to support humans, not supplant them. As we witnessed during the COVID pandemic, safety teams and systems often evolve in unpredictable ways, so technologies and solutions need to be agile to respond to dynamic factors quickly and efficiently.

Data-driven insights

Working collaboratively, regulators and industry are reshaping the safety landscape, specifically the direction of data flow and format of emerging safety sources. We're now at a critical impasse and pharmacovigilance systems must be flexible to adapt to the evolving landscape continually.

Based on more than a decade of experience in designing and deploying a tech-enabled safety risk identification system, we see the following key learnings as critical:

- **Synergy:** Technology is being leveraged successfully by the industry to identify safety risks in any source, volume, and language. Artificial Intelligence (in all its forms), automation technology (in all its forms), and humans (in all our forms) are gifted solutions to the safety industry's challenges but yield far more impressive results combined rather than used in isolation.
- **Needle in the haystack:** Data shows us a tremendous amount of effort is required to find a very small percentage of valid individual case safety reports (ICSRs). There is a lot of noise, but there is also relevant safety information that tech can find with little human effort. Technology best serves when it reduces noise, allowing intelligent insights to be brought forward.
- **Human effort is consistently required:** To validate what 'bubbles up' is of relevance to safety and to assist ML/rewriting of rules and for instances where tech can't make sense of the data as easily (such as in embedded URLs or audio files). Technology is remarkably effective at keeping noise at a minimum to allow humans to focus their attention where it adds the most value, and it can be relied upon to take on more of the workload than humans to identify safety risks regardless of the data source.



Conclusion

When looking to deploy a safety risk identification solution to help comply with regulatory requirements, marketing authorization holders have an increasingly large arsenal of tools to choose from, including AI-enabled and machine-learning systems. Many of these fundamental technologies already exist, and they can be layered onto others for specific use-case scenarios to help identify safety risk from multiple complicated data sets.

Our experience shows that tools such as AI and NLP work best when blended specifically for the scenario or the data source rather than deployed in isolation. It's important to combine accessible data in a meaningful way. With social media, for example, companies may need to update their NLP regularly to make sure it caters to slang, emojis, sentiments, and so on. Identifying risk from patient support program data is similar, as companies need to use multiple tools and technologies because the underlying data is unstructured and messy. The point is to use these techniques in coordination to maximize the benefits of each.

Finally, our experience has shown clearly that although automation brings incredible efficiencies to the drug safety reporting process, human effort is required consistently. There will always be a need for humans to make sure there is no AI bias or any other risks that come with leaning too heavily on AI technology. Human expertise is not only critical at the back end to validate data collected through technologies, but it is also critical on the front end as expert guidance is necessary for selecting the right blend of automation.

Updated continually, IQVIA's Vigilance Detect has been helping pharmaceutical companies manage their pharmacovigilance efforts for 12 years. In that time, we've seen the drug safety landscape evolve and adapt to several increasingly powerful technologies such as artificial intelligence and machine learning. Combined with our ability to process data from disparate sources, the depth and breadth of both our expertise and our technologies enable us to help companies select the right solution for their pharmacovigilance challenges.

About the authors



MARIE FLANAGAN, MSC
Offering Management Director,
IQVIA Vigilance Detect

As Offering Management Director of IQVIA's Vigilance Detect (powered by AE Tracker®), Marie's focus is on developing the Detect portfolio and ensuring the offerings meet client and regulatory needs.

Marie joined Quintiles Drug Safety over 16 years ago. In her tenure, Marie has successfully held many leadership positions in operations and project management and leveraged learned skills and competencies across a broad spectrum of pharmacovigilance (PV) services. In 2017, as the Global Head of PV Services Integration, Marie led a team of PV Product Specialists that focused on the strategic expansion efforts of IQVIA's PV department, conducting PV landscape assessments and analyzing emerging trends. Throughout the years, she has supported Global Business Operations, Strategic Pricing, Business Development, PV Operational Management and the integration of Safety Technology and Services.

Marie graduated from University College Cork, Ireland with an MSc in Medical Microbiology. In 2021, she joined the Vigilance Detect team, leveraging her 16+ years of safety experience.



ALISON SLOANE, MSC, DIP STAT.
General Manager,
IQVIA Vigilance Detect

As General Manager of IQVIA's Vigilance Detect (powered by AE Tracker®), Alison's focus is on driving the vision to provide customers with a tech-enabled optimized approach to adverse event and risk detection in structured and unstructured data. Alison joined Quintiles Drug Safety over 20 years ago. Shortly thereafter, she assumed a customer managed secondment to a pharmaceutical company for 15 months in the UK. During this time, Alison gained experience in a wide range of pharmacovigilance tasks from clinical trials to post marketing and on return to Quintiles she expanded her roles in clinical trials, endpoint management, regulatory reporting and line management. Alison's leadership roles included European leadership of the Pharmacovigilance unit (all functions), global leadership of the Clinical Endpoint Validation and Adjudication (CEVA) Department and subsequently global leadership of the Regulatory Reporting Department, including growing the teams, building out processes, and directing the operational, contractual, financial and customer facing aspects of the organization.

Alison graduated from Trinity College, Dublin, Ireland, with an Honors degree in Natural Science and an MSc and a Diploma in Statistics. Since 2016, Alison's roles have been focused on PV automation and innovation, leveraging her 20+ years of safety experience.

References

1. Arlett, P., Straus, S. and Rasi, G. (2020), Pharmacovigilance 2030. Clin. Pharmacol. Ther., 107: 89-91.
<https://doi.org/10.1002/cpt.1689>
2. CRM Data Based on 0.3M records 2020-2021, 0.23M addressed by Detect ontologies alone.
3. Chatbot/Virtual AI Agent Data based on 0.4M records 2020-2021, 0.29M addressed by Detect ontologies.
4. Case study based on anonymised Detect PSP data current as of Q2 2022: 44,170 digital records received per month 4,300 (9.7%) contain potential adverse events and need to be reviewed by humans/90.3% reduction in manual review.
5. Case study based on anonymised Detect social media and digital data 2018-2022.



CONTACT US

2400 Ellis Road, Durham, NC 27703

iqvia.com/technologies