An industry perspective on the use of machine learning in drug and vaccine safety

Jeffery L. Painter[1], Raymond Kassekert[2] and Andrew Bate[3,4]*

[1]GlaxoSmithKline, Global Safety, Durham, NC, United States, [2]GlaxoSmithKline, Global Safety, Upper Providence, PA, United States, [3]GlaxoSmithKline, Global Safety, Brentford, Middlesex, United Kingdom, [4]London School of Hygiene and Tropical Medicine, London, United Kingdom

In recent years there has been growing interest in the use of machine learning across the pharmacovigilance lifecycle to enhance safety monitoring of drugs and vaccines. Here we describe the scope of industry-based research into the use of machine learning for safety purposes. We conducted an examination of the findings from a previously published systematic review; 393 papers sourced from a literature search from 2000−2021 were analyzed and attributed to either industry, academia, or regulatory authorities. Overall, 33 papers verified to be industry contributions were then assigned to one of six categories representing the most frequent PV functions (data ingestion, disease-specific studies, literature review, real world data, signal detection, and social media). RWD and social media comprised 63% (21/33) of the papers, signal detection and data ingestion comprised 18% (6/33) of the papers, while disease-specific studies and literature reviews represented 12% (4/33) and 6% (2/33) of the papers, respectively. Herein we describe the trends and opportunities observed in industry application of machine learning in pharmacovigilance, along with discussing the potential barriers. We conclude that although progress to date has been uneven, industry is very interested in applying machine learning to the pharmacovigilance lifecycle, which it is hoped may ultimately enhance patient safety.

KEYWORDS

pharmacovigilance, machine learning-ML, drug safety, vaccines safety, artificial intelligence

1 Introduction

The vast increase in the volume of safety reporting over the last years was only exacerbated during the global COVID-19 pandemic. The introduction of new vaccines and medicines in response to the pandemic resulted in more than 1.8 million new safety reports. Enabling the safety community to cope with the onslaught of data has led to increased interest in automating pharmacovigilance (PV) activities within the pharmaceutical industry and prompted safety organizations to advance their technological capabilities (Rudolph et al., 2022).

Even before the pandemic, automation and its potential benefits for PV activities were well recognized (Kassekert et al., 2022). Rules based automation, also known as robotic process automation (RPA), are well established and have been routinely used for several years by companies to assist in the processing of individual case safety reports [ICSRs; see (Kassekert et al., 2020) for an example]. Here, we have adapted the RPA benefit diagram first described by (Vitharanage et al., 2020). Benefits realized through RPA include increased timeliness of data ingestion, higher quality case processing, and reduction in the manual efforts required of case processors. Furthermore, through automation, additional, indirect benefits to PV functions have been realized, such as increased operational reliability, improved job satisfaction, and increased transparency in data management processes (Figure 1). Thematic groups of the nodes are identified by their color, while interdependencies between graph nodes are deduced
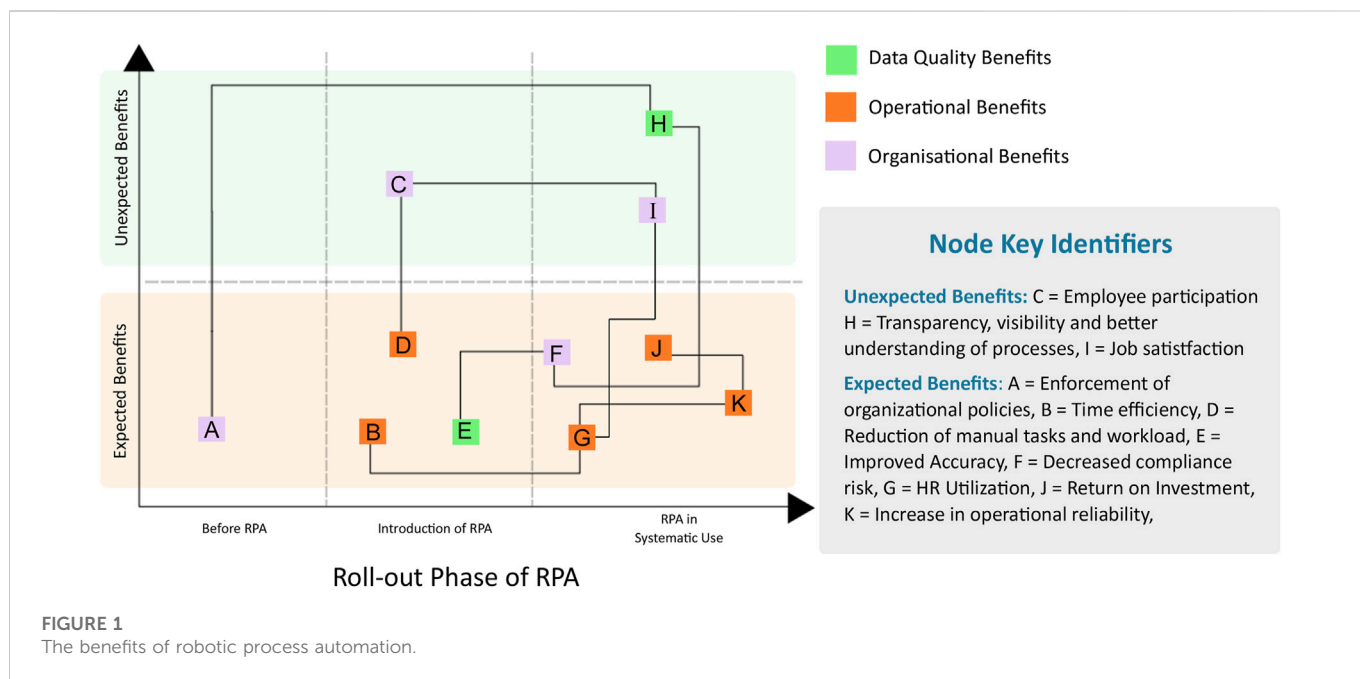
FIGURE 1
The benefits of robotic process automation.

TABLE 1 Details on the benefits of robotic process automation (RPA).

Node	Phase	Benefit	Description
		Expected Benefits	
A	Before RPA	Enforcement of organizational policies	Automation defaults to observing organizational policies in a systematic way
B	Introduction of RPA	Improvement in time efficiency	Reduction in time for case processing and booking
D	Introduction of RPA	Reduction of manual tasks and workload	Reduction of repetitive, mundane, tedious manual tasks
E	Introduction of RPA	Improvement in data accuracy	Automate data processing in systematic way. Decrease potential errors and mistakes in safety case review and processing
F	RPA in Systematic Use	Decrease in compliance risk	Timelier fulfillment of compliance and audit requirements for regulatory bodies
G	RPA in Systematic Use	Improvement in human resource utilization	Better utilization of staff, focus on cases which require human intervention
J	RPA in Systematic Use	Return on investment	Increase capability for high volume case processing without increasing staff
K	RPA in Systematic Use	Increase in operational reliability	Allows the organization to operate reliably, even when faced with unexpected challenges
		Unexpected Benefits	
C	Introduction of RPA	Improved employee participation	Staff is relieved from repetitive tasks and able to focus on improving overall business processes and decision making
H	RPA in Systematic Use	Improvement in transparency, visibility and better understanding of processes	Business processes and business rules are judiciously executed, processes are documented and explainable to stakeholders
I	RPA in Systematic Use	Improvement in job satisfaction	Empower staff to focus on the most important tasks which lead to better job fulfillment

intuitively and deserve future investigation to refine. Each node of the RPA graph is expanded in greater detail with full descriptions in (Table 1).

There has been a growing interest in applying machine learning (ML) across the entire PV lifecycle, from data ingestion and quality control to regulatory or health authority reporting and signal detection. Currently, it is not obvious to what extent ML can be used for routine, operational PV activities and the practical implications this will have for patient safety. More simplistic ML applications such as association rule analysis, or disproportionality

analysis, have been used routinely for signal detection in most organizations (Bate and Evans, 2009). Spontaneous safety reports are routinely analyzed using this family of methods as part of holistic PV management systems (Almenoff et al., 2007). The use of ML in other PV applications however is not routine, although there are isolated reports of ML use for purposes other than signal detection, for example, detection of duplicate spontaneous reports (Norén et al., 2007; Bate and Luo, 2022).

The objective of this paper is to describe the scope of industry-based research into the use of ML for safety purposes, and in particular PV.

2 Unravelling the routine use of ML in PV

A recently published systematic review on the history and use of ML in PV was conducted by evaluating the published literature from 2000–2021 (Kompa et al., 2022). In this scoping review, 393 papers met the criteria for analysis and were analyzed across several metrics including types of safety data used, types of ML models built, and specific PV subtasks addressed. We conducted a thorough examination of the findings of the systematic review (Kompa et al., 2022) to derive a better understanding of the routine use of ML by the pharmaceutical industry and herein present our findings. First, the contributors and affiliations of the 393 papers were reviewed and each paper was attributed to one of three sources (industry, academia, or regulatory authorities). To be included in the original systematic review, the published literature was limited to English papers including "ML terms related to disproportionality analysis, common to PV research, as well as modern ML techniques (e.g., deep learning)." (Kompa et al., 2022). This analysis is subject to these same limitations, and possibly excludes active ML work in the industry which has not been published in peer reviewed literature.

2.1 Attribution assignment

We created an algorithm to automatically perform the attribution function, by searching the authors and affiliations, conflict of interest statements, and sources of funding (grants or sponsorship) of the 393 papers. A list of the top 80 pharmaceutical companies was compiled to allow us to link a particular paper to one or more companies through keyword identification. Regulatory papers were identified similarly using a smaller subset of terms (e.g., FDA, EMA) while the remaining articles were labeled academic papers which did not contain either an industry partner name or regulatory affiliation in either the articles authorship or funding disclosures. The algorithmic generated results for industry-associated publications were then manually reviewed and verified.

Initially, industry-associated articles in this analysis include papers that were either sponsored by one of the companies on the list or have at least one author affiliated with one of the companies on the list. However, in practice, for these industry-associated publications, a broad inclusive definition of 'industry' was taken based on text written in the publications. Publications were included where pharmaceutical industry were involved directly, as co-authors or according to acknowledgements, or sources of funding (i.e., had funded or partially funded the works). Additionally, if more than one academic author declared

industry funding, even if not directly related to the manuscript, we considered it in scope as the foundational work funded by industry could have influenced the work presented in the article. Moreover, as the list of publications was generated for this manuscript using an automated objective filter of the original systematic review, manual review of the output was then used to remove articles not fulfilling our definition of 'industry' for this manuscript from the analysis set.

Forty-three papers were attributed to industry by the algorithmic process, and after manual review of these papers, 33 were verified to be industry contributions. Most of the misclassifications were due to partial term matches of company names. There were less than five articles where the paper did include a company author, but the work was largely regulatory sponsored, and other misses included product mentions of a pharmaceutical company that was not leading the investigation of the manuscript in question. For each of the classifications, all papers were assigned to only one of the three categories (industry, academic, or regulatory).

The breakdown by contributor type provides insight as to the datasets used, the types of models constructed, and the specific PV tasks being addressed by industry, academia, and regulatory bodies (Figure 2A–C).

Many papers used traditional disproportionality analysis (DPA) methods (Figure 3), and such studies made up most of the signal detection research papers across all contributor types.

2.2 Industry trend analysis

When looking at the total number of publications devoted to ML in PV by year (Figure 4), academia dominates, followed by regulatory bodies (e.g., the US Food and Drug Administration, World Health Organization, and Centers for Disease Control and Prevention), whereas industry makes up a small percentage of the overall publications in the field (8.4%, 33/393).

Over the period analyzed (2000–2021), the number of papers attributed to industry, academia, and regulatory authorities has generally increased (Figure 4). Industry appears to have lagged behind academia in the early 2000s, although some articles that could be attributed to industry were not included in the review (Fram et al., 2003; DuMouchel et al., 2004), and it must be acknowledged that some work in ML originating from industry predates the time period included of the review (Alvager et al., 1994; Bate et al., 1998). It may be speculated that industry is not prioritizing publishing these types of papers in the way academia is, and therefore industry may be underrepresented.

Of the 33 industry-associated papers included in the analysis, 19 58% (19/33 = 57.6) were attributed to one of nine different companies, and 14 (42%) were collaborative works that included authors from more than one company.

2.3 Major trends and opportunities in industry application of ML in PV

During the manual review process, each of the 33 industry-associated papers was assigned to one of six primary categories. These were defined by the authors to capture the major PV functions performed by ML in each study. These categories
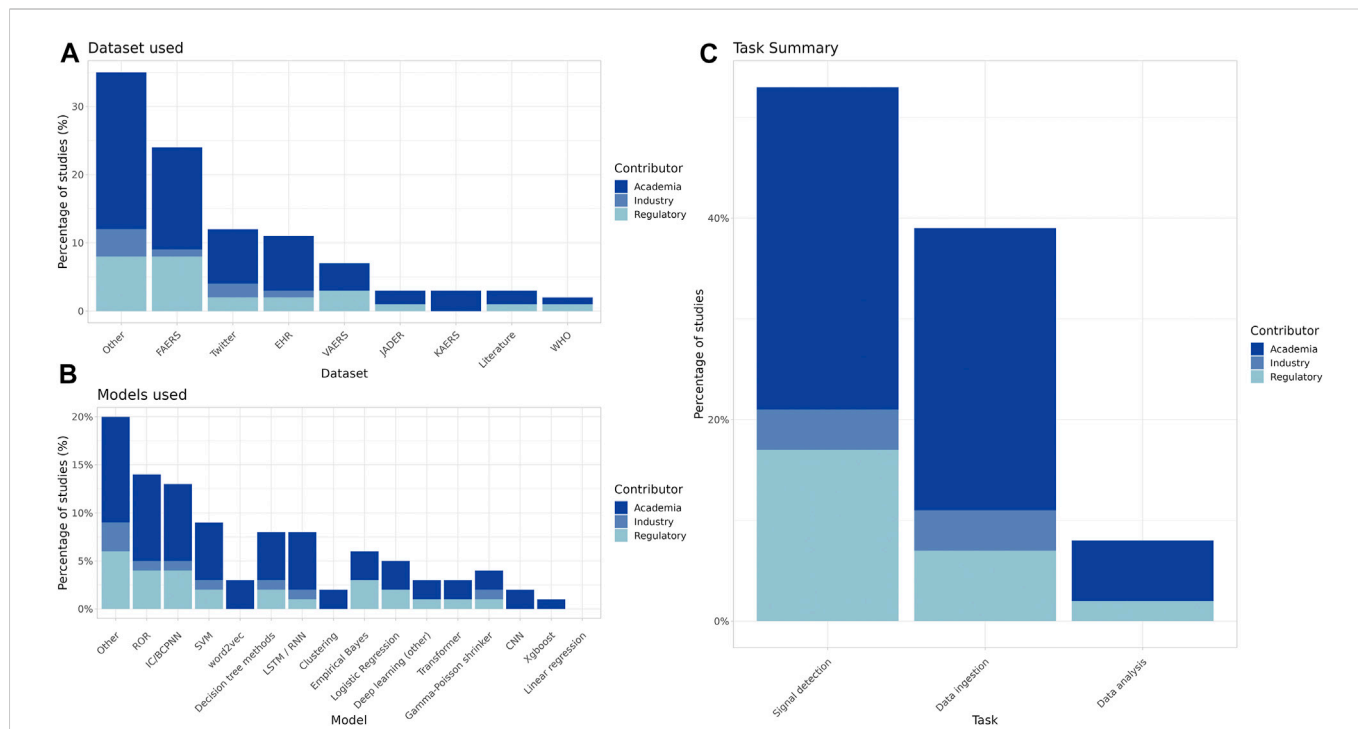
FIGURE 2
Summary of datasets used, primary algorithms, and task type of the included studies by contributor (**(A)**: FAERS = FDA Adverse Event Reporting System, EHR = Electronic Health Records, VAERS = Vaccine Adverse Event Reporting System, JADER = Japanese Adverse Event Reporting System, KAERS = Korean Adverse Event Reporting System, WHO = World Health Organization VigiBase **(B)**: ROR = Reporting Odds Ratio, IC/BCPNN = Information Component/ Bayesian Confidence Propagation Neural Network, SVM = Support Vector Machine, LSTM/RNN = Long short-term memory/Recurrent Neural Network, CNN = Convolutional Neural Network). **(C)** The most common PV tasks for processing and evaluating safety reports include at a high level data ingestion, data analysis and signal detection.
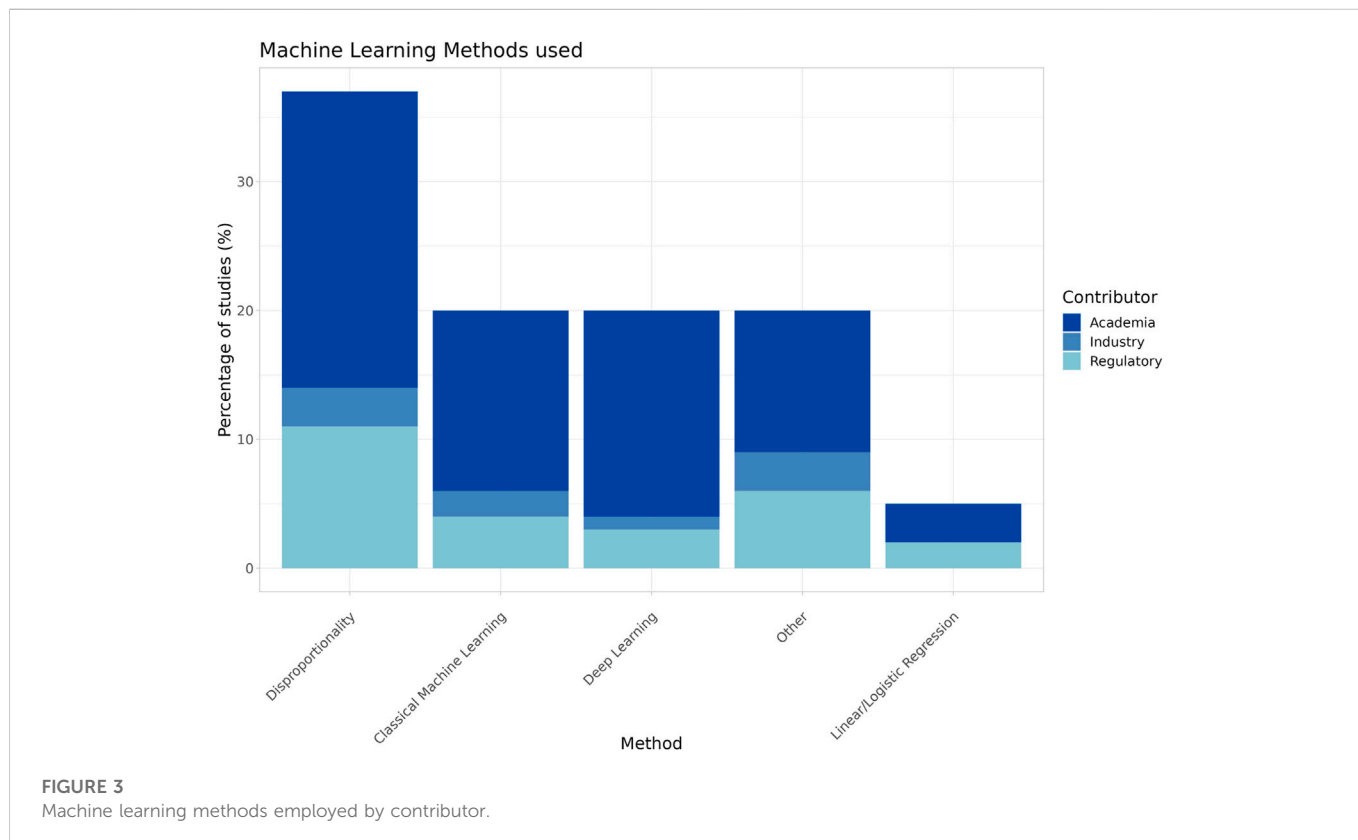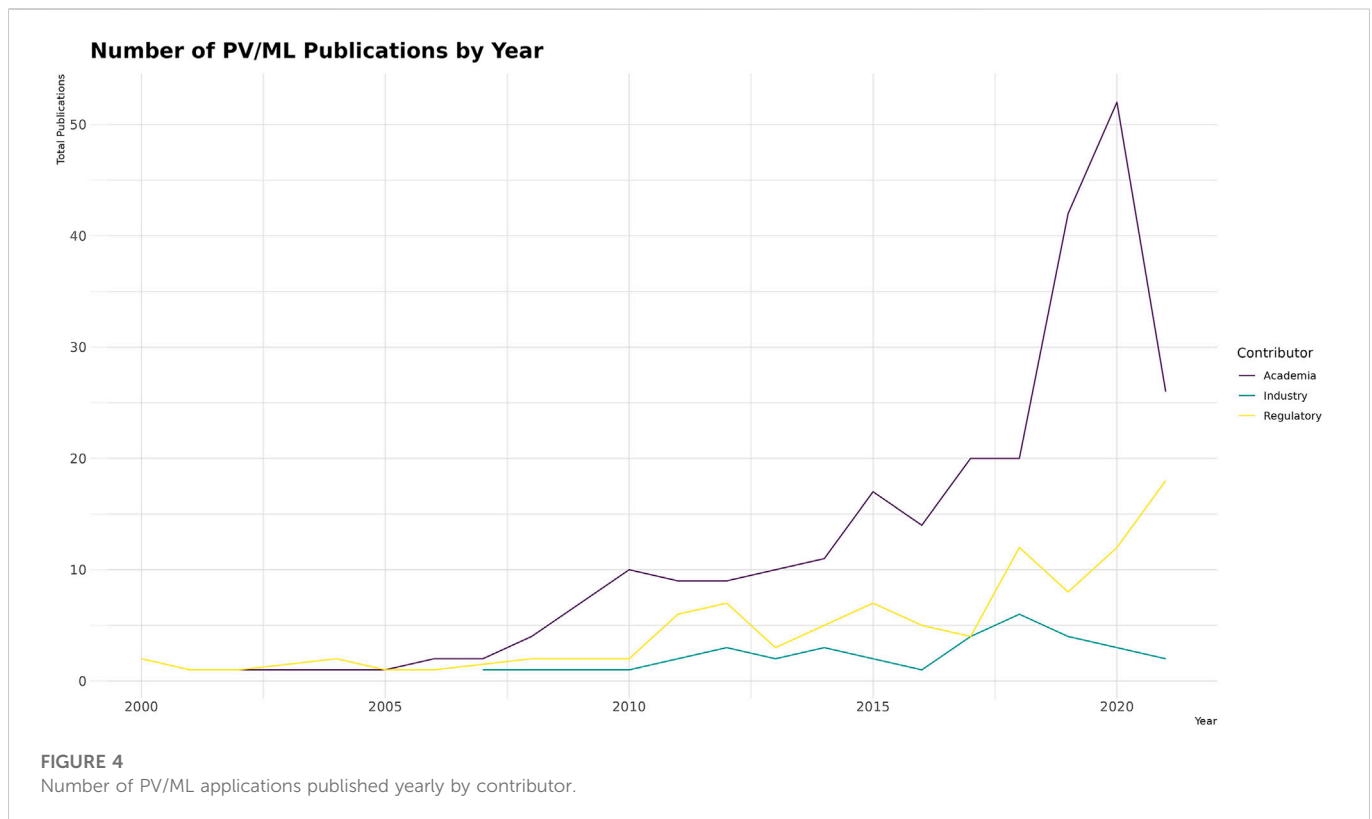

FIGURE 3
Machine learning methods employed by contributor.

FIGURE 4
Number of PV/ML applications published yearly by contributor.

included data ingestion, disease-specific studies, literature review (i.e., for signal detection), leveraging real world data (RWD), signal detection in spontaneous safety reports, and the use of social media data (Figure 5). Topic assignments to each paper were performed manually and can be found in the supplementary data.

RWD (Nordstrom et al., 2007; Gurulingappa et al., 2012a; Cao et al., 2013; Cheetham et al., 2014; Ferrajolo et al., 2014; Yeleswarapu et al., 2014; Whalen et al., 2018; Chapman et al., 2019; Choudhury et al., 2019; Wintzell et al., 2020; Fralick et al., 2021) and social media (Jimeno-Yepes et al., 2015; Powell et al., 2016; Cocos et al., 2017; Curtis et al., 2017; Pierce et al., 2017; Comfort et al., 2018; Gupta et al., 2018; Masino et al., 2018; Gavrielov-Yusim et al., 2019; Gartland et al., 2021) were the most frequent PV functions represented and, collectively, comprised 63% (21/33) of the papers included in the review. This is not surprising, as the use of social media to supplement PV activities began around 2009, with broader use of social media by the public, and attracted interest from both industry and academia as a potential source of safety-related events in near real time. However, more recently, interest in social media data use has declined because of accumulating evidence of variable data quality which limits its value for much (but not all) PV (van Stekelenborg et al., 2019; Powell et al., 2022). In addition, most social media data is typically characterized as unstructured data requiring ML-based methods in order to attempt to glean any insights related to PV activities (Comfort et al., 2018).
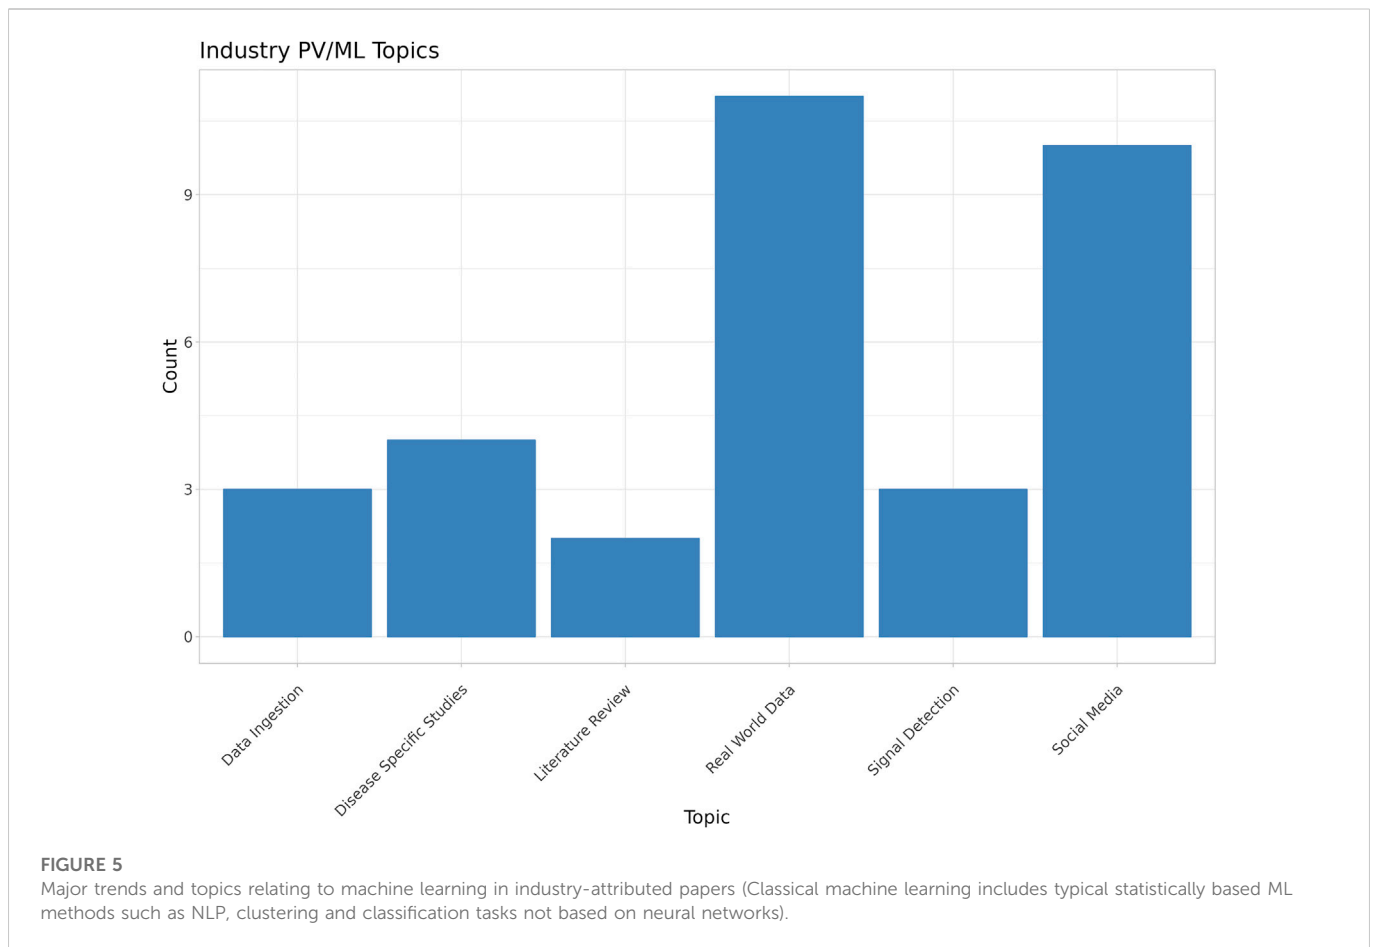
RWD has long been used by industry for pharmacoepidemiologic studies and across the drug and vaccine development lifecycle (Bate et al., 2016; Gatto et al., 2019; Garcia-Gancedo and Bate, 2022). The use of ML for the wider use of RWD continues to attract interest for routine PV activities, and we anticipate this trend to continue. Like social media data, RWD contains unstructured data, and recent research seeks to unlock value from such data in addition to the

structured data in electronic healthcare records. Natural language processing (NLP) and more sophisticated ML is particularly being applied to knowledge extraction from clinical notes in electronic healthcare records (Weiss et al., 2018).

Signal detection in spontaneous reports (Voss et al., 2017; Peng et al., 2020) and data ingestion (Gurulingappa et al., 2013; Abatemarco et al., 2018; Schmider et al., 2019; Routray et al., 2020) represent just under one-fifth (18%, 6/33) of the industry-associated papers included in the review. This aligns with our expectation that routine PV operations utilize ML to automate and improve capabilities. Literature review, which is a routine PV activity, is represented by approximately 6% (2/33) of the papers (Gurulingappa et al., 2012b; Christensson et al., 2012). The remaining studies (12%, 4/33) focused on individual, disease-specific safety issues (Yang et al., 2009; Ratcliffe et al., 2010; Suzuki et al., 2015; Antonazzo et al., 2018).

When compared with the pharmaceutical industry in general, and PV functions in particular, the use of ML appears to be further advanced in other sectors and/or industries, such as manufacturing, finance, and air transportation (Kenyon, 2021; TrifectaDirectory.com, 2021). Given the complexity of medicine, it is difficult, if not impossible, to capture the relevant information in rules (Schwartz et al., 1987) and, given the complexity of PV (Ghosh et al., 2020; Lewis and McCallum, 2020) coupled with large volumes of data, it seems logical that automation and ML will eventually be used routinely and widely in this sector.

Our own experience, and the results of our analysis of the scoping review by Kompa et al. (2022) show that industry is, at this time, very interested in applying ML to the PV lifecycle. As ML tools improve, we expect that they will demonstrate their value above and beyond traditional software and procedural approaches to PV functions, as

FIGURE 5
Major trends and topics relating to machine learning in industry-attributed papers (Classical machine learning includes typical statistically based ML methods such as NLP, clustering and classification tasks not based on neural networks).

is already occurring for routine use of specific tasks, such as employing NLP for more effective screening of literature articles for identifying mentions of suspected adverse events (Glaser et al., 2021). Confidence in, and alignment for the need for such tools is necessary to demonstrate value and engender trust, particularly from regulators. We anticipate that improvements in ML will produce clear benefits to the PV system and may enhance patient safety.

The myriad of overlapping guidance documentation provided to industry causes many to ponder what industry should do, especially if regulatory authorities have different views on the need for and capabilities of ML-enabled PV functions. Clearly fair and safe systems that are trusted by all stakeholders are needed. However, attempting to satisfy all stakeholders may result in a loss of efficiency and the work required to maintain a multifaceted ML-based PV system might exceed the value of the system itself. It is hoped that public-private partnerships such as the Council for International Organizations of Medical Sciences (CIOMS) (Tsintis and La Mache, 2004) may be useful in plotting a course. Indeed, CIOMS recently launched an initiative on AI (Working Group XIV Artificial Intelligence, 2023).

There are many barriers to widespread adoption of ML in PV. These barriers include: the heterogenous nature of PV data; difficulty interpreting the output of ML algorithms; and the lack of performance criteria to determine the acceptability of the output of ML algorithms (Bate and Hobbiger, 2021). Recently, Kassekert et al. (2022) argued that the two major challenges for industry in implementing ML for PV are the risks associated with obtaining adequate training data sets and perceived risk in an emerging regulatory environment (Supplementary Material).

For ML-based systems to succeed in PV, they need to be highly efficient, capable of handling rapid changes in volumes of safety reports and able to learn and incorporate human-in-the-loop mechanisms for identifying novel or unusual patterns and activity. Effective ML should be capable of distinguishing exceptions for human review which may be indicative of data quality issues or a new emerging safety issue (Kjoersvik and Bate, 2022). As this technology evolves, it is important to consider best practices for adoption and validation of these systems, along with consistency of their approach (Huysentruyt et al., 2021).

Case studies describing the application of advanced analytics to perform specific tasks in pharmacovigilance are provided in Table 2.

It is instructive to see how ML is used for safety purposes in other industries. In the air transportation industry, updating accident models has historically been a cumbersome process because of the long period required to review and digest the information contained in long, detailed and highly technical accident reports prepared by safety specialists (Morais et al., 2019). Morais and colleagues developed an ML tool that uses text recognition and text classification, combined with a support vector machine for classifying text according to a predefined taxonomy, to create a 'virtual risk expert' that automatically extracts relevant information from accident reports. The Bayesian network tool was trained on several previous accident reports, while the report for the 2018 Lion Air Boeing 737-8 Max accident provided an opportunity to show the feasibility of the tool for rapidly updating an existing accident model. When the 'virtual

TABLE 2 Case studies describing the application of advanced analytics to perform specific tasks in pharmacovigilance.

Example	Task	Described capability
Identification of a safety outcome through combining NLP extraction of clinical notes with structured data Walker et al. (2016); Weiss et al. (2018)	Computer-assisted elicitation of rules for case definition from experts, using candidate cases drawn from medical records Walker et al. (2016)	The iterative ML-supported technique produced a case definition that had a sensitivity of 92% and a positive predictive value of 79% compared with clinical review alone
	Assess the added contribution of unstructured data extracted from medical text by NLP for detecting acute liver dysfunction (ALD) in patients with inflammatory bowel disease (IBD) Weiss et al. (2018)	Inclusion of NLP terms identified an additional 9% of ALD-onset dates, with consequent earlier recognition in 5% of cases
Social media listening (Twitter and Facebook) for routine post-marketing safety surveillance Powell et al. (2016)	Process 24 months of data to standardize drug names and vernacular symptoms, and to remove duplicates and noise	The methodology effectively transformed social listening data into a usable format for routine post-marketing safety surveillance
Predict the likelihood of a causal association of an observed drug–reaction combination in Individual Case Causality Assessments (ICSR) Cherkas et al. (2022)	Develop a machine learning-based model	Results show that robust probabilistic modeling of ICSR causality is feasible Predictability of the causality assessment of drug–event pairs compared with clinical judgment using global introspection (AUC 0.924; 95% confidence interval [CI] 0.922–0.927). Sensitivity was 0.900 (95% CI 0.896–0.904), and the PPV was 0.778 (95% CI 0.773–0.783)
Signal validation supported by a machine learning (ML)-based pre-validation step to improve process efficiency and consistency Imran et al. (2022)	Cumulative data for six medicinal products [historic signals of disproportionate reporting (SDR) validations and individual case safety reports were used to train and test a ML model	Prediction accuracy of the model ranged from 83% to 86% over 3 months systematic predictions provided valuable information and assisted safety experts in reviewing the SDRs efficiently and consistently
Attitudes of drug safety professionals towards AI in pharmacovigilance Danysz et al. (2019)	Survey of the general sentiments, expectations and readiness for AI	Survey results suggest that pharmacovigilance professionals wish to use their qualifications, skillsets, and experience in work that provides more value for their efforts. Machine learning algorithms have the potential to enhance DS professionals' decision-making processes and support more efficient and accurate case processing

risk expert' was trained exclusively with aviation safety reports it achieved 85% accuracy, whereas if chemical safety reports were included, it achieved 91% accuracy, showing the value of cross-discipline knowledge transfer.

3 Discussion

There has been a marked increase in the use of ML to perform PV functions in the pharmaceutical industry, not just for signal detection, but across the PV lifecycle. The supplementary value of ML when combined with rules-based approaches for fields as broad and complex as medicine, and therefore implicitly PV, makes the inevitability of more widespread ML clear (Schwartz et al., 1987; Rajkomar et al., 2019). As ML use in PV matures, we anticipate seeing even more research, of higher quality and with a greater impact that will eventually lead to routine use of this technology.

While the opportunities are clear, challenges remain to more widespread use of ML across the entire PV lifecycle.

In their scoping review, Kompa et al. (2022) highlighted attributes considered to be best practices in the ML literature. These include appropriate inductive biases, no obvious test-train leakage, tuning hyperparameters, and cross validation. In practical terms, this means using a pre-trained model, rather than building a bespoke model 'from scratch', using external information or data, and using data or code that is in the public domain. Among the 393 papers analyzed, 42 (approximately 10%) reflected these best practices. Of note, most studies (73%) reported using 'off-the-shelf' methods with little to no problem-specific adaptation or domain knowledge. In our analysis, we identified just one industry-

associated paper that clearly reflected these modern best practices. It is important to note that no systematic analysis is exhaustive and inevitably some publications will be missed due to the prespecified criteria limiting the search. Furthermore, not all ML used routinely in industry will be published in papers so there may be some omissions in our perspective resulting from that. Also, as our review focused on publications from a list of the top 80 pharmaceutical companies, we acknowledge there may well be publications from companies outside of this list which are therefore not included in our review. Despite this, we would add that we are unaware of any evidence that suggests any differences in routine ML usage, more broadly and unpublished.

There are many reasons why the use of ML for PV does not more frequently meet or exceed best practice criteria. Whether data even exist or are available is a general challenge in PV, considering for example, the situation in low- and middle-income countries. The sheer volume of data to sift through is often a challenge, especially given the heterogeneity of safety data and issues. Limited access to databases, for privacy reasons or other concerns, is another constraint. These challenges contribute to the discordance between what has been done and what needs to be done to realize the potential of ML in PV.

The TransCelerate Intelligent Automation Opportunities (IAO) and Advancing Safety Analytics (ASA) Initiatives are dedicated to evaluating proposed best practices for the application of interrogative methods to safety data sources (TransCelerate_Biopharma_Inc, 2022). The ASA has issued a white paper that surveyed the current state of signal management, provided a simplified framework that comprises three stages (detection, evaluation and action) and identified best practices (Wisniewski

et al., 2020). More recently, ASA initiative members evaluated the extent of redundancy among three adverse event reporting databases [EudraVigilance Data Analysis System (EVDAS), FDA Adverse Event Reporting System (FAERS) and WHO-VigiBase] by determining the presence or absence of signals of disproportionate reporting (SDRs) for 100 selected products. There were no significant differences in the number and types of safety signals detectable in the three databases, which suggests that each database on its own could be used for signal detection purposes (Vogel et al., 2020). More recently the group has quantified the extent of ICSR replication in terms of the same report being sent to the same recipient (van Stekelenborg et al., 2023). More innovative and effective ways of sharing information could be envisaged and to ensure maximal impact of ML-enabled PV, fundamental changes in PV would be needed (Bate and Stegmann, 2021).

The pharmaceutical industry is investing in automation to perform PV functions with ML. Industry possesses the process expertise and can help identify the business needs for the use of ML, but it will require high technology companies with deep ML knowledge to provide subject matter expertise. Progress on the technology side has been slow to accommodate PV functions, and it can be difficult at this point to separate marketing messages from tangible, demonstrable benefits with respect to proposed software solutions (Hauben et al., 2007). Thus, even as technology improves, manual processing of ICSRs will be required for the foreseeable future.

In conclusion technology holds increasing potential for automating PV functions in the pharmaceutical industry. ML-enabled systems hold great promise. To date, progress has been uneven but there are successes. As barriers to development and implementation can be reduced or resolved, routine use of ML to perform PV functions is likely.

Data availability statement

The original contributions presented in the study are included in the article/Supplementary Material, further inquiries can be directed to the corresponding author.

Supplementary material

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/fdsfr.2023.1110498/full#supplementary-material

References

Abatemarco, D., Perera, S., Bao, S. H., Desai, S., Assuncao, B., Tetarenko, N., et al. (2018). Training augmented intelligent capabilities for pharmacovigilance: Applying deep-learning approaches to individual case safety report processing. *Pharm. Med.* 32, 391–401. doi:10.1007/s40290-018-0251-9

Almenoff, J. S., Powell, G., Schaaf, R., Fram, D., Fitzpatrick, J. M., Pendleton, A., et al. (2007). Online signal management: A systems-based approach that delivers new analytical capabilities and operational efficiency to the practice of pharmacovigilance. *Drug Inf. J. Drug Inf. Assoc.* 41, 779–789. doi:10.1177/009286150704100610

Alvager, T., Smith, T. J., and Vijai, F. (1994). The use of artificial neural networks in biomedical technologies: An introduction. *Biomed. Instrum. Technol.* 28, 315–322.

Antonazzo, I. C., Raschi, E., Forcesi, E., Riise, T., Bjornevik, K., Baldin, E., et al. (2018). Multiple sclerosis as an adverse drug reaction: Clues from the FDA adverse event reporting system. *Expert Opin. Drug Saf.* 17, 869–874. doi:10.1080/14740338.2018.1506763

Bate, A., and Evans, S. J. (2009). Quantitative signal detection using spontaneous ADR reporting. *Pharmacoepidemiol Drug Saf.* 18, 427–436. doi:10.1002/pds.1742

Bate, A., and Hobbiger, S. F. (2021). Artificial intelligence, real-world automation and the safety of medicines. *Drug Saf.* 44, 125–132. doi:10.1007/s40264-020-01001-7

Bate, A., Juniper, J., Lawton, A. M., and Thwaites, R. M. (2016). Designing and incorporating a real world data approach to international drug development and use: What the UK offers. *Drug Discov. Today* 21, 400–405. doi:10.1016/j.drudis.2015.12.002

Bate, A., Lindquist, M., Edwards, I. R., Olsson, S., Orre, R., Lansner, A., et al. (1998). A Bayesian neural network method for adverse drug reaction signal generation. *Eur. J. Clin. Pharmacol.* 54, 315–321. doi:10.1007/s002280050466

Bate, A., and Luo, Y. (2022). Artificial intelligence and machine learning for safe medicines. *Drug Saf.* 45, 403–405. doi:10.1007/s40264-022-01177-0

Bate, A., and Stegmann, J. U. (2021). Safety of medicines and vaccines - building next generation capability. *Trends Pharmacol. Sci.* 42, 1051–1063. doi:10.1016/j.tips.2021.09.007

Cao, H., Lavange, L. M., Heyse, J. F., Mast, T. C., and Kosorok, M. R. (2013). Medical records-based postmarketing safety evaluation of rare events with uncertain status. *J. Biopharm. Stat.* 23, 201–212. doi:10.1080/10543406.2013.735783

Chapman, A. B., Peterson, K. S., Alba, P. R., Duvall, S. L., and Patterson, O. V. (2019). Detecting adverse drug events with rapidly trained classification models. *Drug Saf.* 42, 147–156. doi:10.1007/s40264-018-0763-y

Cheetham, T. C., Lee, J., Hunt, C. M., Niu, F., Reisinger, S., Murray, R., et al. (2014). An automated causality assessment algorithm to detect drug-induced liver injury in electronic medical record data. *Pharmacoepidemiol Drug Saf.* 23, 601–608. doi:10.1002/pds.3531

Cherkas, Y., Ide, J., and Van Stekelenborg, J. (2022). Leveraging machine learning to facilitate individual case causality assessment of adverse drug reactions. *Drug Saf.* 45, 571–582. doi:10.1007/s40264-022-01163-6

Choudhury, O., Park, Y., Salonidis, T., Gkoulalas-Divanis, A., Sylla, I., and Das, A. K. (2019). Predicting adverse drug reactions on distributed health data using federated learning. *AMIA Annu. Symp. Proc.* 2019, 313–322.

Christensson, C., Gipson, G., Thomas, T., and Weatherall, J. (2012). Text analytics for surveillance (TAS):An interactive environment for safety literature review. *Drug Inf. J.* 46, 115–123. doi:10.1177/0092861511428890

Cocos, A., Fiks, A. G., and Masino, A. J. (2017). Deep learning for pharmacovigilance: Recurrent neural network architectures for labeling adverse drug reactions in twitter posts. *J. Am. Med. Inf. Assoc.* 24, 813–821. doi:10.1093/jamia/ocw180

Comfort, S., Perera, S., Hudson, Z., Dorrell, D., Meireis, S., Nagarajan, M., et al. (2018). Sorting through the safety data Haystack: Using machine learning to identify individual case safety reports in social-digital media. *Drug Saf.* 41, 579–590. doi:10.1007/s40264-018-0641-7

Curtis, J. R., Chen, L., Higginbotham, P., Nowell, W. B., Gal-Levy, R., Willig, J., et al. (2017). Social media for arthritis-related comparative effectiveness and safety research and the impact of direct-to-consumer advertising. *Arthritis Res. Ther.* 19, 48. doi:10.1186/s13075-017-1251-y

Danysz, K., Cicirello, S., Mingle, E., Assuncao, B., Tetarenko, N., Mockute, R., et al. (2019). Artificial intelligence and the future of the drug safety professional. *Drug Saf.* 42, 491–497. doi:10.1007/s40264-018-0746-z

Dumouchel, W., Smith, E. T., Beasley, R., Nelson, H., Yang, X., Fram, D., et al. (2004). Association of asthma therapy and churg-strauss syndrome: An analysis of postmarketing surveillance data. *Clin. Ther.* 26, 1092–1104. doi:10.1016/s0149-2918(04)90181-6

Ferrajolo, C., Coloma, P. M., Verhamme, K. M., Schuemie, M. J., De Bie, S., Gini, R., et al. (2014). Signal detection of potentially drug-induced acute liver injury in children using a multi-country healthcare database network. *Drug Saf.* 37, 99–108. doi:10.1007/s40264-013-0132-9

Fralick, M., Kulldorff, M., Redelmeier, D., Wang, S. V., Vine, S., Schneeweiss, S., et al. (2021). A novel data mining application to detect safety signals for newly approved medications in routine care of patients with diabetes. *Endocrinol. Diabetes Metab.* 4, e00237. doi:10.1002/edm2.237

Fram, D. M., Almenoff, J. S., and Dumouchel, W. (2003). "Empirical Bayesian data mining for discovering patterns in post-marketing drug safety," in *Proceedings of the ninth ACM SIGKDD international conference on knowledge discovery and data mining*, Washington, DC, August 24-27, 2003.

Garcia-Gancedo, L., and Bate, A. (2022). Digital biomarkers for post-licensure safety monitoring. *Drug Discov. Today* 27, 103354. doi:10.1016/j.drudis.2022.103354

Gartland, A., Bate, A., Painter, J. L., Casperson, T. A., and Powell, G. E. (2021). Developing crowdsourced training data sets for pharmacovigilance intelligent automation. *Drug Saf.* 44, 373–382. doi:10.1007/s40264-020-01028-w

Gatto, N. M., Sobel, R. E., Geier, J., Mo, J., Bate, A., and Reynolds, R. F. (2019). "The role of pharmacoepidemiology in industry," in *Pharmacoepidemiology*. Editors B. L. Strom, S. E. Kimmel, and S. Hennessy (John Wiley & Sons Ltd), 98–125.

Gavrielov-Yusim, N., Kürzinger, M. L., Nishikawa, C., Pan, C., Pouget, J., Epstein, L. B., et al. (2019). Comparison of text processing methods in social media-based signal detection. *Pharmacoepidemiol Drug Saf.* 28, 1309–1317. doi:10.1002/pds.4857

Ghosh, R., Kempf, D., Pufko, A., Barrios Martinez, L. F., Davis, C. M., and Sethi, S. (2020). Automation opportunities in pharmacovigilance: An industry survey. *Pharm. Med.* 34, 7–18. doi:10.1007/s40290-019-00320-0

Glaser, M., Cranfield, C., Dsouza, D., Duma, A., Hastie, K., Kassekert, R., et al. (2021). Automating individual case safety report identification within scientific literature using natural language processing. *Pharmacoepidemiol. Drug Saf.* 30. Hoboken, NJ: Wiley.

Gupta, S., Pawar, S., Ramrakhiyani, N., Palshikar, G. K., and Varma, V. (2018). Semi-supervised recurrent neural network for adverse drug reaction mention extraction. *BMC Bioinforma.* 19, 212. doi:10.1186/s12859-018-2192-4

Gurulingappa, H., Mateen-Rajput, A., and Toldo, L. (2012a). Extraction of potential adverse drug events from medical case reports. *J. Biomed. Semant.* 3, 15. doi:10.1186/2041-1480-3-15

Gurulingappa, H., Rajput, A. M., Roberts, A., Fluck, J., Hofmann-Apitius, M., and Toldo, L. (2012b). Development of a benchmark corpus to support the automatic extraction of drug-related adverse effects from medical case reports. *J. Biomed. Inf.* 45, 885–892. doi:10.1016/j.jbi.2012.04.008

Gurulingappa, H., Toldo, L., Rajput, A. M., Kors, J. A., Taweel, A., and Tayrouz, Y. (2013). Automatic detection of adverse events to predict drug label changes using text and data mining techniques. *Pharmacoepidemiol Drug Saf.* 22, 1189–1194. doi:10.1002/pds.3493

Hauben, M., Reich, L., Gerrits, C. M., and Younus, M. (2007). Illusions of objectivity and a recommendation for reporting data mining results. *Eur. J. Clin. Pharmacol.* 63, 517–521. doi:10.1007/s00228-007-0279-3

Huysentruyt, K., Kjoersvik, O., Dobracki, P., Savage, E., Mishalov, E., Cherry, M., et al. (2021). Validating intelligent automation systems in pharmacovigilance: Insights from good manufacturing practices. *Drug Saf.* 44, 261–272. doi:10.1007/s40264-020-01030-2

Imran, M., Bhatti, A., King, D. M., Lerch, M., Dietrich, J., Doron, G., et al. (2022). Supervised machine learning-based decision support for signal validation classification. *Drug Saf.* 45, 583–596. doi:10.1007/s40264-022-01159-2

Jimeno-Yepes, A., Mackinlay, A., Han, B., and Chen, Q. (2015). Identifying diseases, drugs, and Symptoms in twitter. *Stud. Health Technol. Inf.* 216, 643–647. doi:10.3233/978-1-61499-564-7-643

Kassekert, R., Easwar, M., Glaser, M., Ventham, R., and Bate, A. (2020). PNS271 automation in routine use for data collection and processing for scalable faster RWE generation. *Value Health* 23, S686. doi:10.1016/j.jval.2020.08.1715

Kassekert, R., Grabowski, N., Lorenz, D., Schaffer, C., Kempf, D., Roy, P., et al. (2022). Industry perspective on artificial intelligence/machine learning in pharmacovigilance. *Drug Saf.* 45, 439–448. doi:10.1007/s40264-022-01164-5

Kenyon, T. (2021). Top 10 sectors for machine learning. Available at: https://aimagazine.com/top10/top-10-sectors-machine-learning.

Kjoersvik, O., and Bate, A. (2022). Black swan events and intelligent automation for routine safety surveillance. *Drug Saf.* 45, 419–427. doi:10.1007/s40264-022-01169-0

Kompa, B., Hakim, J. B., Palepu, A., Kompa, K. G., Smith, M., Bain, P. A., et al. (2022). Artificial intelligence based on machine learning in pharmacovigilance: A scoping review. *Drug Saf.* 45, 477–491. doi:10.1007/s40264-022-01176-1

Lewis, D. J., and Mccallum, J. F. (2020). Utilizing advanced technologies to augment pharmacovigilance systems: Challenges and opportunities. *Ther. Innov. Regul. Sci.* 54, 888–899. doi:10.1007/s43441-019-00023-3

Masino, A. J., Forsyth, D., and Fiks, A. G. (2018). Detecting adverse drug reactions on twitter with convolutional neural networks and Word embedding Features. *J. Healthc. Inf. Res.* 2, 25–43. doi:10.1007/s41666-018-0018-9

Morais, C., Yung, K., and Patelli, E. (2019). "Machine-learning tool for human factors evaluation - application to Lion Air Boeing 737-8 Max accident," in *UNCECOMP 2019. 3rd ECCOMAS thematic conference on international conference on uncertainty quantification in computational Sciences and engineering*. Editors M. Papadrakakis, V. Papadopoulos, and G. Stefanou (Greece: UNCECOMP), 24.

Nordstrom, B. L., Norman, H. S., Dube, T. J., Wilcox, M. A., and Walker, A. M. (2007). Identification of abacavir hypersensitivity reaction in health care claims data. *Pharmacoepidemiol Drug Saf.* 16, 289–296. doi:10.1002/pds.1337

Norén, G. N., Orre, R., Bate, A., and Edwards, I. R. (2007). Duplicate detection in adverse drug reaction surveillance. *Data Min. Knowl. Discov.* 14, 305–328. doi:10.1007/s10618-006-0052-8

Peng, L., Xiao, K., Ottaviani, S., Stebbing, J., and Wang, Y. J. (2020). A real-world disproportionality analysis of FDA Adverse Event Reporting System (FAERS) events for baricitinib. *Expert Opin. Drug Saf.* 19, 1505–1511. doi:10.1080/14740338.2020.1799975

Pierce, C. E., Bouri, K., Pamer, C., Proestel, S., Rodriguez, H. W., Van Le, H., et al. (2017). Evaluation of Facebook and twitter monitoring to detect safety signals for medical products: An analysis of recent FDA safety alerts. *Drug Saf.* 40, 317–331. doi:10.1007/s40264-016-0491-0

Powell, G. E., Seifert, H. A., Reblin, T., Burstein, P. J., Blowers, J., Menius, J. A., et al. (2016). Social media listening for routine post-marketing safety surveillance. *Drug Saf.* 39, 443–454. doi:10.1007/s40264-015-0385-6

Powell, G., Kara, V., Painter, J. L., Schifano, L., Merico, E., and Bate, A. (2022). Engaging patients via online healthcare fora: Three pharmacovigilance use cases. *Front. Pharmacol.* 13, 901355. doi:10.3389/fphar.2022.901355

Rajkomar, A., Dean, J., and Kohane, I. (2019). Machine learning in medicine. *N. Engl. J. Med.* 380, 1347–1358. doi:10.1056/NEJMra1814259

Ratcliffe, S., Younus, M., Hauben, M., and Reich, L. (2010). Antidepressants that inhibit neuronal norepinephrine reuptake are not associated with increased spontaneous reporting of cardiomyopathy. *J. Psychopharmacol.* 24, 503–511. doi:10.1177/0269881108100776

Routray, R., Tetarenko, N., Abu-Assal, C., Mockute, R., Assuncao, B., Chen, H., et al. (2020). Application of augmented intelligence for pharmacovigilance case seriousness determination. *Drug Saf.* 43, 57–66. doi:10.1007/s40264-019-00869-4

Rudolph, A., Mitchell, J., Barrett, J., Sköld, H., Taavola, H., Erlanson, N., et al. (2022). Global safety monitoring of COVID-19 vaccines: How pharmacovigilance rose to the challenge. *Ther. Adv. Drug Saf.* 13, 20420986221118972. doi:10.1177/20420986221118972

Schmider, J., Kumar, K., Laforest, C., Swankoski, B., Naim, K., and Caubel, P. M. (2019). Innovation in pharmacovigilance: Use of artificial intelligence in adverse event case processing. *Clin. Pharmacol. Ther.* 105, 954–961. doi:10.1002/cpt.1255

Schwartz, W. B., Patil, R. S., and Szolovits, P. (1987). Artificial intelligence in medicine. Where do we stand? *N. Engl. J. Med.* 316, 685–688. doi:10.1056/NEJM198703123161109

Suzuki, A., Yuen, N. A., Ilic, K., Miller, R. T., Reese, M. J., Brown, H. R., et al. (2015). Comedications alter drug-induced liver injury reporting frequency: Data mining in the WHO VigiBase. *Regul. Toxicol. Pharmacol.* 72, 481–490. doi:10.1016/j.yrtph.2015.05.004

Transcelerate_Biopharma_Inc (2022). Advancing safety analytics solutions. Available at: https://www.transceleratebiopharmainc.com/assets/advancing-safety-analytics-solutions/ (Accessed November 6, 2022).

Trifectadirectory.Com (2021). 6 Sectors embracing AI & ML technology. Available at: https://www.trifectadirectory.com/6-sectors-embracing-ai-ml-technology/.

Tsintis, P., and La Mache, E. (2004). CIOMS and ICH initiatives in pharmacovigilance and risk management: Overview and implications. *Drug Saf.* 27, 509–517. doi:10.2165/00002018-200427080-00004

Van Stekelenborg, J., Ellenius, J., Maskell, S., Bergvall, T., Caster, O., Dasgupta, N., et al. (2019). Recommendations for the use of social media in pharmacovigilance: Lessons from IMI WEB-RADR. *Drug Saf.* 42, 1393–1407. doi:10.1007/s40264-019-00858-7

Van Stekelenborg, J., Kara, V., Haack, R., Vogel, U., Garg, A., Krupp, M., et al. (2023). Individual case safety report replication: An analysis of case reporting transmission networks running head: Replication of case safety reports. *Drug Saf.* 46, 39–52. doi:10.1007/s40264-022-01251-7

Vitharanage, I. D., Bandara, W., Syed, R., and Toman, D. (2020). "An empirically supported conceptualisation of robotic process automation (RPA) benefits," in *the 28th European Conference on Information Systems (ECIS 2020)*.

Vogel, U., Van Stekelenborg, J., Dreyfus, B., Garg, A., Habib, M., Hosain, R., et al. (2020). Investigating overlap in signals from EVDAS, FAERS, and VigiBase. *Drug Saf.* 43, 351, doi:10.1007/s40264-019-00899-y

Voss, E. A., Boyce, R. D., Ryan, P. B., Van Der Lei, J., Rijnbeek, P. R., and Schuemie, M. J. (2017). Accuracy of an automated knowledge base for identifying drug adverse reactions. *J. Biomed. Inf.* 66, 72–81. doi:10.1016/j.jbi.2016.12.005

Walker, A. M., Zhou, X., Ananthakrishnan, A. N., Weiss, L. S., Shen, R., Sobel, R. E., et al. (2016). Computer-assisted expert case definition in electronic health records. *Int. J. Med. Inf.* 86, 62–70. doi:10.1016/j.ijmedinf.2015.10.005

Weiss, L. S., Zhou, X., Walker, A. M., Ananthakrishnan, A. N., Shen, R., Sobel, R. E., et al. (2018). A case study of the incremental utility for disease identification of natural language processing in electronic medical records. *Pharm. Med.* 32, 31–37. doi:10.1007/s40290-017-0216-4

Whalen, E., Hauben, M., and Bate, A. (2018). Time series disturbance detection for hypothesis-free signal detection in longitudinal observational databases. *Drug Saf.* 41, 565–577. doi:10.1007/s40264-018-0640-8

Wintzell, V., Svanström, H., Melbye, M., Ludvigsson, J. F., Pasternak, B., and Kulldorff, M. (2020). Data mining for adverse events of tumor Necrosis factor-alpha Inhibitors in pediatric patients: Tree-based scan statistic Analyses of Danish Nationwide health data. *Clin. Drug Investig.* 40, 1147–1154. doi:10.1007/s40261-020-00977-5

Wisniewski, A., Gomez, A., Jokinen, J., Lacroix, K., Garg, A., Grabowski, N., et al. (2020). *Signal management: Current landscape and considerations for best practices*. TransCelerate Biopharma Inc. Available: http://www.transceleratebiopharmainc.com/wp-content/uploads/2020/04/TransCelerate_ASA_SignalManagementManuscript_April2020.pdf (Accessed November 6, 2022).

Working Group XIV Artificial Intelligence (2023). Council for international organizations of medical Sciences. Available at: https://cioms.ch/working_groups/working-group-xiv-artificial-intelligence-in-pharmacovigilance/.

Yang, X., Brandenburg, N. A., Freeman, J., Salomon, M. L., Zeldis, J. B., Knight, R. D., et al. (2009). Venous thromboembolism in myelodysplastic syndrome patients receiving lenalidomide: Results from postmarketing surveillance and data mining techniques. *Clin. Drug Investig.* 29, 161–171. doi:10.2165/00044011-200929030-00003

Yeleswarapu, S., Rao, A., Joseph, T., Saipradeep, V. G., and Srinivasan, R. (2014). A pipeline to extract drug-adverse event pairs from multiple data sources. *BMC Med. Inf. Decis. Mak.* 14, 13. doi:10.1186/1472-6947-14-13