



Black Swan Events and Intelligent Automation for Routine Safety Surveillance

Oeystein Kjoersvik¹ · Andrew Bate^{2,3,4}

Accepted: 27 February 2022
© TransCelerate BioPharma, Inc. 2022

Abstract

Effective identification of previously implausible safety signals is a core component of successful pharmacovigilance. Timely, reliable, and efficient data ingestion and related processing are critical to this. The term ‘black swan events’ was coined by Taleb to describe events with three attributes: unpredictability, severe and widespread consequences, and retrospective bias. These rare events are not well understood at their emergence but are often rationalized in retrospect as predictable. Pharmacovigilance strives to rapidly respond to potential black swan events associated with medicine or vaccine use. Machine learning (ML) is increasingly being explored in data ingestion tasks. In contrast to rule-based automation approaches, ML can use historical data (i.e., ‘training data’) to effectively predict emerging data patterns and support effective data intake, processing, and organisation. At first sight, this reliance on previous data might be considered a limitation when building ML models for effective data ingestion in systems that look to focus on the identification of potential black swan events. We argue that, first, some apparent black swan events—although unexpected medically—will exhibit data attributes similar to those of other safety data and not prove algorithmically unpredictable, and, second, standard and emerging ML approaches can still be robust to such data outliers with proper awareness and consideration in ML system design and with the incorporation of specific mitigatory and support strategies. We argue that effective approaches to managing data on potential black swan events are essential for trust and outline several strategies to address data on potential black swan events during data ingestion.

1 Introduction

The ability to effectively capture, organise, compile, analyse, share, and act on highly heterogenous data in a timely manner is at the heart of pharmacovigilance [1]. As data volumes grow [2], coupled with the complexity and variability of international regulation requirements, there is a pressing need to ensure that emerging scientific and technological advances are harnessed to ensure data management remains workable and effective and that the focus remains on patient

Key Points

‘Black swan events’ are unexpected, severe situations that, in retrospect, can seem predictable. A subset of safety risks seen in pharmacovigilance are such black swan events.

Adequate data ingestion is core to safety. Machine learning (ML) during pharmacovigilance data ingestion should make both data intake and processing more effective and enable signal detection and management, or—at an absolute minimum—not hinder it.

Routine use of ML in pharmacovigilance data ingestion requires considering the potential for black swan events. It needs to support both adequate ingestion of familiar and common reporting patterns or attributes and unexpected changes in reporting that might signify a black swan event.

There are many manifestations of potential black swan events in pharmacovigilance data, but ML can be anticipated to support or enable the adequate ingestion of data on these events if ML best practice and—when needed—mitigatory strategies are employed.

✉ Andrew Bate
andrew.x.bate@gsk.com

¹ R&D IT, MSD, Prague, Czech Republic

² Global Safety, GSK, 980 Great West Road, Brentford TW8 9GS, Middlesex, UK

³ Department of Non-Communicable Disease Epidemiology, London School of Hygiene and Tropical Medicine, London, UK

⁴ Department of Medicine at NYU Grossman School of Medicine, New York, USA

safety [3]. Pharmacovigilance systems are necessarily complex as they ingest data from very different sources. They are therefore made up of many distinct process steps with different purposes across data intake, processing, and reporting. All of this must happen in a timely and effective manner [4].

‘Intelligent automation’ uses automation technologies to streamline and scale decision making across organizations. Intelligent automation simplifies processes, frees up resources, improves operational efficiencies, and has various applications [5]. Intelligent automation incorporates robotic process automation and ML with other processes to enable rapid and effective intake, processing, and reporting of heterogeneous data at scale. Of course, it is paramount that systems incorporating algorithmic learning can identify and appropriately manage unexpected, and even rare, data points to build trust in these systems.

2 Trusted Intelligent Automation and Data Ingestion

Given the large amount of manual effort routinely required in pharmacovigilance operations to effectively manage safety data, the promise of intelligent automation in pharmacovigilance is enormous [6, 7]. Rule-based systems are in production use, for example, in duplicate checking and data quality review [8, 9]. More and more opportunities for robotic process automation will be identified, and solutions and increasing numbers of solutions will therefore continue to have a clear impact on the automation of routine tasks. However, the complexity and breadth of medical knowledge limits the possibility of capturing all relevant information in rules [10]. The additional benefits of automation will be seen as ML capabilities in safety evolve to supplement rule-based data ingestion and communication [6]. ML has advanced enormously in recent years, and we believe that what are now standard ML approaches from other industries and areas can and will be increasingly applied in pharmacovigilance. Although the usefulness of ML will be limited for some current tasks by the current volume of training data [6], for other tasks—particularly those that are the same or similar to those in broader healthcare—no such limitation should be anticipated, e.g., voice input technology to facilitate work in electronic healthcare records [11]. Ultimately, ML will be routinely used across all parts of the pharmacovigilance data lifecycle, from data ingestion to analysis and augmented decision making, as automation becomes increasingly ‘intelligent’.

Effectively negotiating the regulatory challenges of validating emerging ML-based technologies is also key to building confidence and trust among pharmacovigilance stakeholders [12].

ML credibility and reliability are based on demonstrated learnings from the most seen types of training data ingested

in development. Although this may be representative and voluminous, it does not represent the universe of potential inputs. Further trust can be built by appropriately identifying and learning from these outliers by, for example, computing a low confidence score or triaging for manual (human) assessment.

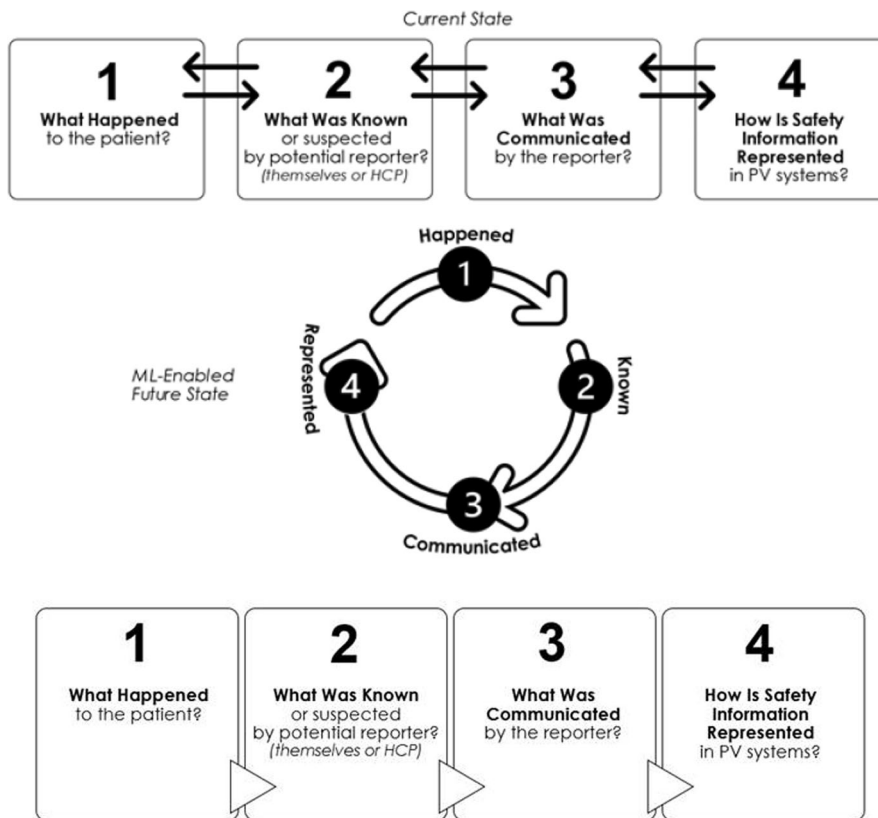
Intelligent automation solutions must ultimately show that a combination of rules-based systems and ML can ensure, as effectively and reliably as possible, that data in safety databases are promptly available for analysis and reflect all known and relevant information from a suspected adverse event (AE) occurrence (see Fig. 1).

The use of ML in safety can be broadly distinguished by asking two questions [6]. First, can ML further our ability to find novel data patterns as we conduct signal detection or consider data aspects in our evaluations that might otherwise be suppressed or hidden and therefore missed (e.g., variable combinations for propensity score estimation in pharmacoepidemiology studies) [13]? Second, to what extent can we do existing data ingestion work, particularly the more routine operationally repetitive tasks, more efficiently and effectively through learning from data? Or take on new work that was not previously possible [4]? For example, ‘autoencoding’ verbatim text to *International Conference on Harmonisation Medical Dictionary for Regulatory Activities* (MedDRA) code AE terms during data entry [14] is an ML opportunity, assuming sufficient training data exist to enable ML to systematically handle the data, where pre-identified specific rules may be impossible or at least limiting. Although the first use is out of scope for this manuscript, we should note that the analytic challenges to resolve across these two pharmacovigilance application areas are often very similar (e.g., identifying medical events in free text) [15]. However, the aims of ML in data entry and data output analysis are very different: ML for data analysis on the outputs is to enable the ability to find interesting patterns and outliers; ML in automation focuses on effective learning from previous data to adapt processes and approaches more effectively. As an example, ML in data ingestion might help filter out noise and errors while ensuring that the original core data are represented with excellent fidelity and minimal information loss.

2.1 The Concept of Black Swan Events

Although unexpected findings affect all elements of drug (and vaccine) development, this is perhaps most keenly felt in the field of safety surveillance, where—in addition to surveillance of known or suspected safety issues and the potential occurrence in routine healthcare practice—the identification of previously unanticipated safety issues is of paramount importance. In refining our understanding of the safety of routinely

Fig. 1 ICSR process overview and how a ICSR ML enabled process could approach the challenge of preventing loss of fidelity of data. *HCP* healthcare practitioner, *ML* machine learning, *PV* pharmacovigilance



used medicines, the capability to effectively identify previously unexpected safety signals is crucial [16].

Even a rare safety issue can lead to drastic changes in the recommended usage or even the withdrawal of a medicine or vaccine already used in routine healthcare delivery. An accurate understanding of the emerging safety profiles of medicines and vaccines during development is vital for effective and appropriate progressing of drug and vaccine development programmes.

The concept of ‘black swan events’ is of relevance to pharmacovigilance. Taleb [17] defines a black swan event as one of extreme impact that, although outside the realm of regular expectations (i.e., prospectively unpredictable), prompts humans to concoct explanations for its occurrence after the fact, making it seemingly explainable and predictable (i.e., retrospectively distorted). Humans are often irrational when considering risk and probabilities [18]. Sandman et al. [19] argued that the public’s response to risks is driven by the nature of the hazard and their outrage that such events could happen, accompanied by a strong urge to blame someone if and when they do occur. This tendency naturally leads to retrospective distortion as selective attention is given to different sources of risk. The social amplification of risk can be seen as an ex-post attempt to make up for failures to anticipate extreme events [20]. In addition to impact and retrospective predictability, all black swans are ‘outliers’ as they are unexpected given what has been seen before.

Europeans referred to a black swan as something impossible until actual black swans were discovered in Australia in 1697. The term black swan event has been perhaps most used in the financial industry, for example, to refer to the global financial crisis of 2008–2009 [21]. It has also been used to refer to the Ebola epidemic in Africa in 2015 [22]. Most recently, the COVID-19 pandemic has been referred to as a black swan event in terms of its widespread impact [23, 24].

Taleb’s definition of a black swan event as prospectively unpredictable depends on the observer. An event that was a black swan to one person (or algorithm) may have been predictable and therefore not a black swan to another. Presumably, at least some Australians in 1697 would have dismissed the notion of white swans. One needs to be aware of the possible existence of future potential black swans, and therefore systems and processes should detect, or at least flag during ingestion, data that might signify potential black swan events proactively and prospectively.

3 Black Swan Events in Safety Data

The concept of black swan events has only received cursory reference in the pharmacovigilance literature, and not in the context of intelligent automation [25, 26]. We consider a pharmacovigilance black swan event to be a drug or

vaccine–AE combination that became a new, unexpected safety signal and then had a big impact on the benefit–risk profile of the medicine/vaccine, which some may have subsequently felt could have been identified earlier with the benefit of hindsight. For this manuscript, we define ‘big impact’ as a safety risk that significantly affects the known benefit–risk profile of a medicine or vaccine, leading to changes in medicine or vaccine utilization. ‘Unexpected’ could mean that data accruing are systematically different from what was previously in the pharmacovigilance system (data in general or, more narrowly, increased reporting of a specific drug/vaccine–AE pair; i.e., a ‘data outlier’) and/or that the emerging risk was unpredictable in terms of contemporaneous medical knowledge, irrespective of reporting pattern (a ‘medical outlier’). Table 1 shows examples of different historic pharmacovigilance safety signals and/or risks that could be considered black swan events. They are distinct in terms of the nature of medical knowledge at the time of risk identification or in terms of data outlier manifestation. Black swan events may be either medical or data outliers or both data and medical outliers. This table emphasizes the variety of potential black swan events that a pharmacovigilance system must look to identify a priori.

ML at its core requires learning from data, and algorithmic performance relies on sound, contemporaneous, up-to-date training data to learn. The use of ML in automated ingestion raises concern that outliers, including potential black swan events, may be smoothed away into the generality of the data as noise rather than attracting the focus they deserve as data points that may be systematically different from the other observed data and therefore require further investigation. Black swan events present one specific challenge to the trusted use of ML approaches in pharmacovigilance. ML researchers actively study the handling of ‘out-of-distribution’ data and providing uncertainty estimates for model predictions [27–30].

Much of the safety data to be ingested into pharmacovigilance systems are neither unexpected (e.g., reporting of well-established safety issues) nor serious (mild), and some have minimal potential to impact the understood benefit–risk profile of a medicine [3]. A credible ML-enabled intelligent automation process for data ingestion needs to appropriately ingest the data manifestation of these known/familiar/predictable AEs and appropriately handle the rare new information, including potential pharmacovigilance black swan events. The data manifestation of a potential black swan event would be data outliers in quantitative terms, i.e., drug–event pairs more frequently reported than might have been anticipated, or aspects of their reports or clinical outliers that are clinically convincing data on safety reports alluding to clinical novelty (see Table 1).

4 Machine Learning to Pre-Empt, Alert, and Mitigate Potential Black Swans During Safety Data Ingestion

In Sect. 4.1, we list some ML implementation strategies to pre-empt and alert to potential black swan events during data ingestion, so that signal detection with any approach is at least as likely to—as quickly as possible—identify the emerging safety issue as would have been possible with more manual data ingestion processes. Table 2 shows a range of examples of how ML might be applied to aspects of data ingestion and illustrates the pitfalls of injudicious ML implementation that could lead to poor representation of data in safety databases, potentially impacting signal detection. On the other hand, there is no certainty, of course, that manual curation and processing have always optimally helped in the identification of black swan events, and one can anticipate that, in the future, ML perhaps might identify some such deficiencies.

Further perspective is needed here: As discussed earlier, black swan events are considered from a specific human vantage point. Safety is littered with examples that some would consider black swan events, although some of these safety issues may have seemed unexpected from a contemporaneous medical knowledge viewpoint while exhibiting more typical patterns of accumulating data seen with many emerging safety issues. Such typical data patterns would be increasing volumes of reports describing a specific vaccine/drug–event pair or related concepts, etc., and the increasingly informative cumulative data conveyed on reports represents typical pharmacovigilance information, e.g., well-known coded adverse drug reaction terms and other information important in safety data assessment, so that the ‘unexpectedness’ is focussed on the lack of plausibility between the drug/vaccine and AE. Such safety issues would therefore present in such a way that effective intelligent automation with ML would be unlikely to be problematic from an algorithmic perspective in terms of case intake and processing, enabling the data to be efficiently handled with minimum human involvement while retaining the fidelity of the stored and transmitted data to that originally received. Some of the black swan event examples in Table 1 have data patterns similar to those seen previously (volume and nature of data conveyed on safety reports); some do not. Therefore, emerging issues with different data patterns require different mitigatory surveillance strategies. There may be a lack of data for some, or data may differ from what was seen previously. Although we concur with others that there is enormous potential for more effective learning from data for medicines and vaccines [31], practical approaches are needed to handle and learn effectively from ‘predictable’ data (i.e., data similar

Table 1 Examples of different drug or vaccine–adverse event pairs that could be considered pharmacovigilance black swan events and how these exemplars can inform future machine learning data ingestion strategies

Black swan event	Safety reporting of exemplar	Safety reporting?	Data outlier? ^a	Medical outlier ^b	Implications for future ML use for data ingestion
RotaShield vaccine and intussusception	Emerging safety issue in real-world use with plausibility at the time [51]	Yes	Yes (pair) ^c	No	Only data outlier is unexpectedly frequent reporting of vaccine–AE pairs ML algorithms should work if statistical reporting patterns are similar in character to how previous safety issues were reported as they emerged
Cisapride QT prolongation	QT prolongation is a known frequent ADR to drugs. Safety reports of this drug–AE pair accumulated, triggering clinical review. Contemporaneous medical knowledge and other data made an association seem implausible and refuted by some [52–54]	Yes	Yes (pair) ^c	Yes	Only data outlier is unexpectedly frequent reporting of drug–AE pair combination ML algorithms should work if statistical reporting patterns are similar in character to how previous safety issues were reported as they emerged If the link between a drug/vaccine and an AE is medically implausible, data should be available with high fidelity as reported AE is well known and captured in PV systems
Practolol sclerosing with peritonitis as part of oculo musculo cutaneous syndrome	A novel, previously unknown ADR received non-specific AE reporting [55, 56]	Yes	Yes	Yes	Presents as a data outlier and a medical outlier Monitoring of ML is essential as subtle changes in reporting patterns may be indicative of an emerging new type of issue ^d
Silicone breast implants—autoimmune-like disorders	Delayed or lack of timely reporting into PV systems [57–59]	No	NA	Yes	A data outlier only in the sense of absence of data in the safety system Awareness that ML-based systems cannot mitigate effectively for lack of data Necessitates ability to conduct surveillance in other data streams, particularly without reliance on the necessity of suspicion of an issue

Although out of the scope of this manuscript, further ML deployment in signal detection and analysis might also facilitate a more effective medical review of the safety report (e.g., by case clustering)

ADR adverse drug reaction, AE adverse event, ML machine learning, NA not applicable, PV pharmacovigilance

^aIngested data listed on safety reports at time of identification significantly different to data on reports seen previously

^bUnexpected because contemporaneous medical knowledge at the time of identification seemed to render an association unlikely

^cData outlier solely in terms of unexpected reporting of drug/vaccine–event pair but not in other ways

^dNote that signals of potential subgroup risks, new information on a signal (e.g., occurrence at different dosage), and other local differences in reporting could also present as data outliers—a subset of which could also be potential black swan events

Table 2 Examples of using machine learning in the data ingestion process where inappropriate deployment could complicate downstream identification of an emerging signal

ML task	Purpose of ML task within overall PV data ingestion process	Black swan event potential consequence if ML inappropriately deployed during data ingestion
Privacy-preserving algorithms to ensure that personalized data are not collected and shared further	PV data without PII are made available for analysis [60, 61]	In the removal of PII, additional data that would be important or insightful for PV activity are removed
Identification of similar safety reports through learned unexpectedness of similar attributes (or combinations of) from previous reporting patterns	Clustered unexpectedly similar case reports indicative of higher probability of being duplicate safety reports [62]	Similar reports suddenly reported that are systematically different from those seen before flagged as duplicates of the same instance rather than novel emerging
ML used to find likely data errors due to the identification of unrealistic data outliers and highlight and/or propose adjustments for such data errors	ML used to look for likely erroneous data elements as part of the quality assurance process for safety data; for example, recording an extreme dose and dosage form combination [63]	True, potentially important data outliers would be less visible at the analysis stage or even wrongly adjusted
Extract data of relevance from extensive data input with much extraneous information, e.g., audio reporting of an adverse safety event or a chatbot transcript [64]	ML extracts relevant PV information from communications for safety reports [4, 7]	Critical safety information articulated in original communication not appropriately captured

ML machine learning, PII personal identifiable information, PV pharmacovigilance

in properties to historic data such as ongoing reporting of common well-established AEs) but also to pre-empt and be aware of potential black swan events and optimally learn from data outliers. Such practical approaches could realize the full potential of data-driven technologies in pharmacovigilance and, fortunately, are well-studied in the technical ML literature; see the discussions of ‘out-of-distribution learning’ in Hendrycks et al. [30], Blundell et al. [32], Shafaei et al. [33], and Meinke et al. [34].

4.1 Considerations for Pharmacovigilance Use of Machine Learning to Avoid and Pre-Empt Potential Black Swan Events

The following points should be considered when using ML-based systems to enable pre-emption of and alerting to potential black swan events in data ingestion and related tasks in pharmacovigilance.

- Have in place a pre-planned mitigatory process to provide essential confidence and trust in the use of ML systems. If a black swan does occur, have a process to deal with unexpected outcomes.
- ‘Data drift’ can occur over time as models are used in production [35, 36]. That is, the data can gradually change over time, perhaps almost imperceptibly becoming very distinct from the data at production launch: This property of data drift needs to be monitored by defining and then complying with ‘significant change’ triggers for human review [37, 38]. Such a trigger is a metric or metrics that measure against model performance and, if a threshold is breached, will redirect for a human to review more extensively instead of the routine use of the ML model. Humans are generally better suited to deal with unexpected patterns outside of the expected scope, and significant change triggers will be a pre-defined threshold that forces the system to require careful human review over its trained model. Such periodic reviews may lead to changes to the ML model through retraining on new data and redeploying the updated algorithm.
- Continually monitor for unanticipated changes in model performance metrics (e.g., leveraging extreme value theory, which focuses on statistical properties of rare events) [39].
- Ensure metadata, which explain the data, are kept current for linking, grouping, and interpreting data fields for increased insight into the data profile.
- Orthogonal data mining (i.e., using different algorithms that are as independent as possible to discover previously unknown patterns/outliers representing fragments or indicators of emerging black swans) reduces the probability of missing an outlier by using two distinct methods. Similarly, ensemble methods [27, 40] can also be used;

these create multiple distinct models and then combine them to produce improved results.

- Issue periodic model robustness challenges to find any weaknesses/limiting assumptions, such as the use of ‘adversarial attack’ approaches, which look to deliberately subvert otherwise reliable ML systems [41]; continually assess the impact of input to the model (e.g., with sensitivity analyses and input data perturbation); and observe changes in the amount or distribution of missing data. This approach would aim to understand model behaviour when interacting with unexpected data to create a more robust system or model.
- Model prediction needs to account for rare events that do not occur in the training data; an ML method should always compute and make the level of confidence in predictions transparent to users and always have a way of returning ‘NULL’, ‘prediction unable to be made’, and ‘requires human review’ responses when prediction is not possible [42–44].
- Frequently update the model with new training data to ensure ongoing model generalizability. Model retraining should be triggered by a metric(s) measuring a drift in the data to evaluate model performance. Cost-sensitive loss functions should be applied to judge the quality and capability of the models, which look to weigh up the cost of missing a rare event. If retraining is not triggered within a set amount of time, consider setting a time-based trigger for retraining.
- Articulate clearly any uncertainty in ML models using methodological approaches. For example, if using an neural network model, add a diversity term in the loss function during training [45] and consider the uncertainty to the fullest extent possible in resulting use [46] to make the potential for when and where black swan events might occur as clear as possible to users.
- Enhance credibility and useability to the extent possible by explainability, not only of how the model works and underlying assumptions but also regarding uncertainty in the model. Ensure trust in the evaluation and assessment methods used to develop the ML approach and maximize the reproducibility of ML analyses [47] to enable others to test and challenge the approach and reduce the chance of black swan events.
- Employ existing explanation methods for prediction models to better understand the influential input parameters of a model and thereby enable one to determine whether a model is likely to be robust to changes in the environment or confounding factors [48].
- Appropriate training of the users of the ML application is essential [49].
- When practising ML, the pre-processing step is part of the model and should have precision and recall metrics

associated with it, and the performance of the whole system should ideally be measured or constructed at once for ‘end-to-end learning’, starting from raw data with deep learning.

- Data should never be smoothed or left out to make the ML easier to execute. Instead, ensure enough data, and use methods that address the ambiguity.
- Critical analysis of known and likely mechanistic factors influencing treatment and diseases could be leveraged to better understand what factors of variation a model would have to be robust to and to formalise how changes in our understanding over time may necessitate model adaptation [50].

5 Conclusion

ML is increasingly being explored in pharmacovigilance as a component of pharmacovigilance systems to enable adequate data ingestion and related tasks. Systems for intelligent automation need to effectively manage common and reported safety data and rarer data points such as black swan events. Safety needs to harness the standard approaches used in the broader ML community as well as recent advances and implement ML systems that recognise the possibility of rare events and are designed accordingly. Black swan events are, by their nature, a fact of life. The design of ML-based systems for pharmacovigilance needs to recognize and actively plan for the possibility of such events.

With trusted and reliable strategies to avoid and mitigate potential black swan events, overall systems leveraging ML can realize their maximum impact for data intake and processing, making overall surveillance activities more effective, including for safety.

Acknowledgements The authors gratefully acknowledge Clint Craun of TransCelerate Biopharma Inc. for editing support and Lea Goetz for editing support mainly focussed on the ML elements.

Declarations

Funding Work on intelligent automation that led to the co-authors writing this manuscript was supported by TransCelerate BioPharma Inc.

Conflicts of interest Oeystein Kjoersvik is an employee of MSD. Andrew Bate is an employee and stockholder of GSK.

Ethics approval Not applicable.

Consent to participate Not applicable.

Consent for publication Not applicable.

Availability of data and materials Not applicable.

Code availability Not applicable.

Authors' contributions All authors contributed to the manuscript conception and design. All authors read and approved the final manuscript.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License, which permits any non-commercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc/4.0/>.

References

- Lindquist M. Data quality management in pharmacovigilance. *Drug Saf.* 2004;27(12):857–70.
- Stergiopoulos S, Fehrl M, Caubel P, Tan L, Jebson L. Adverse drug reaction case safety practices in large biopharmaceutical organizations from 2007 to 2017: an industry survey. *Pharmaceut Med.* 2019;33(6):499–510.
- Bate A, Stegmann JU. Safety of medicines and vaccines—building next generation capability. *Trends Pharmacol Sci.* 2021;42(12):1051–63.
- Ghosh R, Kempf D, Pufko A, Barrios Martinez LF, Davis CM, Sethi S. Automation opportunities in pharmacovigilance: an industry survey. *Pharmaceut Med.* 2020;34(1):7–18.
- IBM. What is intelligent automation. 2021. <https://www.ibm.com/cloud/learn/intelligent-automation>
- Bate A, Hobbiger SF. Artificial intelligence, real-world automation and the safety of medicines. *Drug Saf.* 2021;44(2):125–32.
- Lewis DJ, McCallum JF. Utilizing advanced technologies to augment pharmacovigilance systems: challenges and opportunities. *Ther Innov Regul Sci.* 2020;54(4):888–99.
- Kassekert R, Easwar M, Glaser M, Venham R, Bate A. Automation in routine use for data collection and processing for scalable faster RWE generation. *Value Health.* 2020 (**in Press**).
- Glaser M, Cranfield C, Dsouza D, Duma A, Hastie K, Kassekert R, et al. Automating individual case safety report identification within scientific literature using natural language processing. *Pharmacoepidemiol Drug Saf.* 2021;30:118–881.
- Rajkomar A, Dean J, Kohane I. Machine learning in medicine. *N Engl J Med.* 2019;380(14):1347–58.
- Kumah-Crystal YA, Pirtle CJ, Whyte H, Goode ES, Anders SH, Lehmann CU. Electronic health record interactions through voice: a review. *Appl Clin Inform.* 2018;9(03):541–52.
- Huysentruyt K, Kjoersvik O, Dobracki P, Savage E, Mishalov E, Cherry M, et al. Validating intelligent automation systems in pharmacovigilance: insights from good manufacturing practices. *Drug Saf.* 2021;44(3):261–72.
- Sessa M, Khan AR, Liang D, Andersen M, Kulahci M. Artificial intelligence in pharmacoepidemiology: a systematic review. Part 1—overview of knowledge discovery techniques in artificial intelligence. *Front Pharmacol.* 2020;11:1028.
- Brown EG, Wood L, Wood S. The medical dictionary for regulatory activities (MedDRA). *Drug Saf.* 1999;20(2):109–17.
- Létinier L, Jouganous J, Benkebil M, Bel-Létoile A, Goehrs C, Singier A, et al. Artificial intelligence for unstructured health-care data: application to coding of patient reporting of adverse drug reactions. *Clin Pharmacol Therap.* 2021;130:392.
- Bate A, Evans SJ. Quantitative signal detection using spontaneous ADR reporting. *Pharmacoepidemiol Drug Saf.* 2009;18(6):427–36.
- Taleb NN. Black swans and the domains of statistics. *Am Stat.* 2007;61(3):198–200.
- Spiegelhalter D. Risk and uncertainty communication. *Annu Rev Stat Appl.* 2017;4:31–60.
- Sandman PM, Miller PM, Johnson BB, Weinstein ND. Agency communication, community outrage, and perception of risk: three simulation experiments. *Risk Anal.* 1993;13(6):585–98.
- Kasperson RE, Renn O, Slovic P, Brown HS, Emel J, Goble R, et al. The social amplification of risk: a conceptual framework. *Risk Anal.* 1988;8(2):177–87.
- Bekiros S, Boubaker S, Nguyen DK, Uddin GS. Black swan events and safe havens: the role of gold in globally integrated emerging markets. *J Int Money Financ.* 2017;73:317–34.
- Osterholm MT, Moore KA, Gostin LO. Public health in the age of Ebola in West Africa. *JAMA Intern Med.* 2015;175(1):7–8.
- Gray GL, Alles MG. Measuring a business's grit and survivability when faced with “black swan” events like the coronavirus pandemic. *J Emerg Technol Acc.* 2021;18(1):195–204.
- Yarovaya L, Matkovskyy R, Jalan A. The effects of a “black swan” event (COVID-19) on herding behavior in cryptocurrency markets. *J Int Financ Markets Inst Money.* 2021;75:101321.
- Edwards IR. Causality assessment in pharmacovigilance: still a challenge. *Drug Saf.* 2017;40(5):365.
- Fan BE, Shen JY, Lim XR, Tu TM, Chang CCR, Khin HSW, et al. Cerebral venous thrombosis post BNT162b2 mRNA SARS-CoV-2 vaccination: a black swan event. *Am J Hematol.* 2021;96(9):E357–61.
- Lakshminarayanan B, Pritzel A, Blundell C. Simple and scalable predictive uncertainty estimation using deep ensembles. *arXiv preprint arXiv:161201474.* 2016.
- Farina F, Phillips L, Richmond NJ. Intrinsic uncertainties and where to find them. *arXiv preprint arXiv:210702526.* 2021.
- Ovadia Y, Fertig E, Ren J, Nado Z, Sculley D, Nowozin S, et al. Can you trust your model's uncertainty? Evaluating predictive uncertainty under dataset shift. *arXiv preprint arXiv:190602530.* 2019.
- Hendrycks D, Mazeika M, Dietterich T. Deep anomaly detection with outlier exposure. *arXiv preprint arXiv:181204606.* 2018.
- Finelli LA, Narasimhan V. Leading a digital transformation in the pharmaceutical industry: reimagining the way we work in global drug development. *Clin Pharmacol Ther.* 2020;108(4):756–61.
- Blundell C, Cornebise J, Kavukcuoglu K, Wierstra D. Weight uncertainty in neural network. *International Conference on Machine Learning*; 2015: PMLR; 2015. p. 1613–22.
- Shafaei A, Schmidt M, Little JJ. A less biased evaluation of out-of-distribution sample detectors. *arXiv preprint arXiv:180904729.* 2018.
- Meinke A, Bitterwolf J, Hein M. Provably Robust Detection of Out-of-distribution Data (almost) for free. *arXiv preprint arXiv:210604260.* 2021.
- Ditzler G, Roveri M, Alippi C, Polikar R. Learning in non-stationary environments: a survey. *IEEE Comput Intell Mag.* 2015;10(4):12–25.
- Finlayson SG, Subbaswamy A, Singh K, Bowers J, Kupke A, Zittrain J, et al. The clinician and dataset shift in artificial intelligence. *N Engl J Med.* 2021;385(3):283–6.
- Chandra SR. Scalable and secure learning with limited supervision over data streams. <https://utd-ir.tdl.org/bitstream/handle/>

- 10735.1/6196/ETD-5608-011-CHANDRA-8457.95.pdf?sequence=6&isAllowed=y: Texas; 2018.
38. Ackerman S, Farchi E, Raz O, Zalmanovici M, Dube P. Detection of data drift and outliers affecting machine learning model performance over time. arXiv preprint arXiv:201209258. 2020.
 39. Lund R. Revenge of the white swan. *Am Stat*. 2007;61(3):189–92.
 40. Dietterich TG. Ensemble methods in machine learning. International workshop on multiple classifier systems; 2000: Springer; 2000. p. 1–15.
 41. Finlayson SG, Bowers JD, Ito J, Zittrain JL, Beam AL, Kohane IS. Adversarial attacks on medical machine learning. *Science*. 2019;363(6433):1287–9.
 42. Gennatas ED, Friedman JH, Ungar LH, Pirracchio R, Eaton E, Reichmann LG, et al. Expert-augmented machine learning. *Proc Natl Acad Sci*. 2020;117(9):4571–7.
 43. Madras D, Pitassi T, Zemel R. Predict responsibly: improving fairness and accuracy by learning to defer. arXiv preprint arXiv:171106664. 2017.
 44. Mozannar H, Sontag D. Consistent estimators for learning to defer to an expert. In: International Conference on Machine Learning; 2020: PMLR; 2020. p. 7076–87.
 45. Wabartha M, Durand A, Francois-Lavet V, Pineau J. Handling black swan events in deep learning with diversely extrapolated neural networks. Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20, International Joint Conferences on Artificial Intelligence Organization; 2020; 2020. p. 2140–7.
 46. Kompa B, Snoek J, Beam AL. Second opinion needed: communicating uncertainty in medical machine learning. *NPJ Digit Med*. 2021;4(1):1–6.
 47. McDermott MBA, Wang S, Marinsek N, Ranganath R, Foschini L, Ghassemi M. Reproducibility in machine learning for health research: still a ways to go. *Sci Transl Med*. 2021;13(586).
 48. Molnar C. Interpretable machine learning: Lulu. com; 2020.
 49. Danysz K, Cicirello S, Mingle E, Assuncao B, Tetarenko N, Mockett R, et al. Artificial intelligence and the future of the drug safety professional. *Drug Saf*. 2019;42(4):491–7.
 50. Peters J, Janzing D, Schölkopf B. Elements of causal inference: foundations and learning algorithms. The MIT Press; 2017.
 51. Markatou M, Ball R. A pattern discovery framework for adverse event evaluation and inference in spontaneous reporting systems. *Stat Anal Data Mining ASA Data Sci J*. 2014;7(5):352–67.
 52. Olsson S, Edwards IR. Tachycardia during cisapride treatment. *BMJ*. 1992;305(6856):748–9.
 53. Inman W, Kubota K. Tachycardia during cisapride treatment. *BMJ*. 1992;305(6860):1019.
 54. Layton D, Key C, Shakir SA. Prolongation of the QT interval and cardiac arrhythmias associated with cisapride: limitations of the pharmacoepidemiological studies conducted and proposals for the future. *Pharmacoepidemiol Drug Saf*. 2003;12(1):31–40.
 55. Bate A, Lindquist M, Orre R, Edwards IR, Meyboom RH. Data-mining analyses of pharmacovigilance signals in relation to relevant comparison drugs. *Eur J Clin Pharmacol*. 2002;58(7):483–90.
 56. Mann RD. An instructive example of a long-latency adverse drug reaction—sclerosing peritonitis due to practolol. *Pharmacoepidemiol Drug Saf*. 2007;16(11):1211–6.
 57. Brewer T, Colditz GA. Postmarketing surveillance and adverse drug reactions: current perspectives and future needs. *JAMA*. 1999;281(9):824–9.
 58. Kessler DA. Introducing MEDWatch. A new approach to reporting medication and device adverse effects and product problems. *JAMA*. 1993;269(21):2765–8.
 59. Bate A, Reynolds RF, Caubel P. The hope, hype and reality of Big Data for pharmacovigilance. *Ther Adv Drug Saf*. 2018;9(1):5–11.
 60. LePendu P, Iyer SV, Bauer-Mehren A, Harpaz R, Mortensen JM, Podchiyska T, et al. Pharmacovigilance using clinical notes. *Clin Pharmacol Ther*. 2013;93(6):547–55.
 61. Bhattacharya M, Snyder S, Malin M, Truffa MM, Marinic S, Engelmann R, et al. Using social media data in routine pharmacovigilance: a pilot study to identify safety signals and patient perspectives. *Pharmaceut Med*. 2017;31(3):167–74.
 62. Norén GN, Orre R, Bate A, Edwards IR. Duplicate detection in adverse drug reaction surveillance. *Data Min Knowl Disc*. 2007;14(3):305–28.
 63. Star K, Caster O, Bate A, Edwards IR. Dose variations associated with formulations of NSAID prescriptions for children: a descriptive analysis of electronic health records in the UK. *Drug Saf*. 2011;34(4):307–17.
 64. Nath J. Chatbot, machine learning and artificial intelligence in pharmacovigilance: maintaining privacy, optimizing efficiency. 2018 [cited 2021 25th November]; <https://chatbotsmagazine.com/chatbot-machine-learning-and-artificial-intelligence-in-pharmacovigilance-maintaining-privacy-877283e4b4b7>. Accessed 11 Mar 2022.