



# Back to the Future: The Evolution of Pharmacovigilance in the Age of Digital Healthcare

# 22

Michael A. Ibara and Rachel L. Richesson

## Abstract

Pharmacovigilance is the science and activities relating to the detection, assessment, understanding, and prevention of adverse effects from medicines or vaccines. Pharmacovigilance originated in an attempt to better understand the safety of drugs in order to ensure and protect the safety individual patients and consumers. Over time, the development of the field has been heavily influenced by the need for the pharmaceutical industry to fulfill regulatory requirements, with the unintended result of losing track of the individual patient. With the onset of digitized healthcare data, we have an opportunity to reunite the industrial and personal in pharmacovigilance to increase the scope and efficiency of monitoring and the speed of response. Informatics supports this transformation by advancing a pharmacovigilance research agenda that should include defining conceptual (ontological) and operational definitions for adverse events that can address dif-

ferent product types and regulatory contexts, developing standards and systems to detect and report adverse events at scale and from different data sources, and developing methods (including artificial intelligence and machine learning) to predict risks of adverse events an various populations.

## Keywords

Pharmacovigilance · Informatics · Adverse drug events · Postmarketing surveillance  
Pharmacoepidemiology · Quantitative signal detection

## Learning Objectives

1. Define the term pharmacovigilance and describe how pharmacovigilance relates to assuring the safety of medications and vaccines.
2. Define the terms “adverse event,” “adverse drug event,” and “adverse drug reaction,” list the four required elements from a regulatory perspective, and discuss the relationship of an adverse event to the notion of causality to a specified medication or vaccine product.
3. Discuss the relationship between adverse events of regulatory interest to adverse events of medical interest and the emerging role of electronic health record data for each.

---

M. A. Ibara (✉)  
Elligo Health Research, Princeton, NJ, USA  
e-mail: [michael.ibara@elligodirect.com](mailto:michael.ibara@elligodirect.com)

R. L. Richesson  
Department of Learning Health Sciences, University  
of Michigan Medical School, Ann Arbor, MI, USA  
e-mail: [richessr@med.umich.edu](mailto:richessr@med.umich.edu)

4. Discuss the need for a modernized operational definition for adverse events in the current era, and describe how healthcare and pharmacovigilance workflows and systems can change to accommodate use of electronic health records in AE detection, investigation, and reporting.
5. List and describe four specific informatics topics in need of require research and development to support a modern pharmacovigilance infrastructure.

*The ideas which are here expressed so laboriously are extremely simple and should be obvious. The difficulty lies, not in the new ideas, but in escaping from the old ones, which ramify, for those brought up as most of us have been, into every corner of our minds.* (John Maynard Keynes; from the preface to *The General Theory of Employment, Interest, and Money* 1936)

## Introduction

This chapter seeks to provide a foundation for future work in pharmacovigilance for the informatician involved in clinical research. It will not attempt to provide an overview of the field of pharmacovigilance, as this has been covered extensively elsewhere [1] including a previous version of this chapter [2]. Important definitions and resources are presented in Table 22.1. The focus here will be on key developments in pharmacovigilance and related areas as a result of the growing digitization of healthcare data. We will propose an informatics research agenda meant to move the field forward and provide for a more holistic consideration of patient safety.

Pharmacovigilance is defined by the World Health Organization (WHO) [5] as “the science and activities relating to the detection, assessment, understanding and prevention of adverse effects or any other drug-related problem.” Pharmacovigilance is a central practice for understanding and assuring drug safety. For an excellent history of the development of pharmacovigilance as a discipline and the general applicability of informatics, see the previous edition of this chapter [2] in which the authors

provide a superb primer for those wishing to gain a better understanding of the topic. A full treatment of the historical, regulatory, industrial, statistical, and medical aspects of the field can be found in several excellent reference works on the topic, especially *Stephens’ Detection and*

**Table 22.1** Definitions [2–4] and relevant organizations

An <b>adverse event (AE)</b> is broadly defined as any clinical event, sign, or symptom that goes in an unwanted direction. Adverse events also include worsening of preexisting conditions, per FDA and ICH definitions. No assertion of causality is implied with adverse events
An <b>adverse drug event (ADE)</b> is harm caused by appropriate or inappropriate use of a drug
An <b>adverse drug reaction (ADR)</b> includes the suggestion of a causal relationship (e.g., probable, possible) between the event and a therapeutic agent or device. Adverse drug reactions are a subset of ADEs, where harm is directly caused by a drug under appropriate use (i.e., at normal doses). After an ADR is suspected (i.e., adverse consequences are speculated to be caused from a drug), then careful and systematic data collection is required to evaluate that suspicion for further action
The <b>international conference on harmonization (ICH E2B)</b> issues international safety reporting guidance
The <b>US Food and Drug Administration (FDA)</b> supports the FDA adverse event reporting system (FAERS) database of AE and medication error reports and product quality complaints (resulting in AEs) submitted to FDA as part of FDA’s postmarketing safety surveillance program for drug and therapeutic biologic products. Reporting requirements are summarized at: <a href="https://open.fda.gov/data/faers/">https://open.fda.gov/data/faers/</a>
The <b>European medicines agency (EMA)</b> is a decentralized agency of the European Union (EU) responsible for the scientific evaluation, supervision and safety monitoring of medicines in the EU
The <b>Council for International Organizations of medical sciences (CIOMS)</b> has been instrumental in developing pharmacovigilance standards and practice. Both CIOMS and ICH operate as forums for discussion and standardization of drug safety methods and requirements
The <b>WHO international drug monitoring program</b> , supported and coordinated by the WHO collaborating Centre for International Drug Monitoring (“the Uppsala monitoring Centre”), serves as a globally integrated and deliberate pharmacovigilance system and maintains the international database of adverse drug events
Adverse events and medication errors are coded using terms in the <b>medical dictionary for regulatory activities (MedDRA)</b>

*Evaluation of Adverse Drug Reactions: Principles and Practice* [6] and *Mann's Pharmacovigilance* [4].

---

## Background

Pharmacovigilance originated as an attempt to better understand the safety of drugs in order to protect individual patients and improve medicine. But while today pharmacovigilance plays a key and vital role in the research and public health arena, to the uninitiated, it can seem bureaucratic, arcane, and arbitrary. This is due mainly to the myriad influences on the field from medicine, public health, industry, and regulation, as well as from broad interest in the topic by patients and practitioners, academic and industry researchers, and regulatory and legal bodies—all groups who have a stake in the endeavor. Over time, pharmacovigilance has taken on the shape of these combined influences, and their often disparate demands have led to a balkanization of the original pharmacovigilance landscape. Today, what a biopharmaceutical industry professional would describe as the daily work of pharmacovigilance would be unrecognizable to the layman or even to healthcare researchers in safety not otherwise engaged with industry.

For a number of years, the most significant forces of differentiation in pharmacovigilance were (and remain) the regulatory and legal requirements to which drug and device manufacturers must comply (hence the often-quoted statement by industry professionals that “compliance” is their first priority). And while healthcare practitioners are subjected to significant regulations and laws as well, a difference in focus and content means that “drug safety” in a healthcare or academic research setting has come to mean something quite different from the industrial use of the term. As the field developed over the last 50 years, the patient was seen as the recipient of any learning and good practices in research on safety of drugs and medical devices but was only taken seriously as a participant at the level of their individual healthcare provider. Both indus-

trial and academic researchers saw the patient more as a source rather than a collaborator in their own health and well-being.

The result of these trends is that, today, pharmacovigilance looks very much like the rest of healthcare: siloed and having difficulty interoperating with other healthcare components. With separate standards, processes, systems, and data stores, various practitioners of “drug safety” work on their individual agendas, not noticing or acknowledging that they share (or could share) the same data with researchers in other fields of pharmacovigilance. But today, the increasing digitization of healthcare data is challenging this compartmentalization as it becomes possible to have a single data source serve a host of downstream practitioners and researchers, as well as the empowered patient.

What is less obvious, but we argue even more significant, is that the digitization of healthcare data creates the possibility for a return to the original aspirations of the field—where we can recapture the original goals of pharmacovigilance and reunite the individual, population, academic, and industrial pursuits to an extent that benefits all stakeholders but most especially which allows us to realize one of the original goals of pharmacovigilance, to protect the individual while contributing to greater understanding at a population level. Practitioners in academic, medical, and industrial settings are finding themselves more often than not pursuing and working with the same data from the same sources. It is encouraging to imagine that they will also work on research topics that will help to reunify the field of pharmacovigilance and move it forward.

The previous edition of this chapter [7] provided a full exposition of the Coasian economic approach [8, 9] to pharmacovigilance to support the thesis that the digitization of healthcare data creates opportunities to unify the field of pharmacovigilance. As this has been amply demonstrated over the last several years, here we provide a summary of the approach of Coasian economics as it applies to the field of Pharmacovigilance, below.

## Coasian Transactions: The Development and Evolution of PV/Drug Safety

The Coasian development of pharmacovigilance can be outlined as follows:

1. Historically, pharmacovigilance was largely developed by vertical organizations having the resources to find, collect, and process safety information—drug and device manufacturers.
2. These organizations were the de facto owners of safety information and responsible for it (focus of regulations) because they were the only organizations able to afford the transaction costs.
3. As healthcare data has become digitized, there has been a dramatic lowering of the “transaction cost” of finding, collecting, and reporting safety information.
4. The movement of AE transaction costs toward zero means that the economic incentives to maintain vertical organizations for pharmacovigilance will no longer be present.
5. With AE transaction able to be horizontally (across different organizations), this creates an environment where new business models and opportunities are encouraged.

If we view pharmacovigilance through a Coasian lens, we see that not only what we call adverse events but also related healthcare data which may impact our assessments, or which can be used in novel ways to improve our ability to practice pharmacovigilance, will continue to increase in number and at an increasing rate, for the foreseeable future.

The challenge for us is to unify (or reunify) the very different professional guilds that have developed as previously described. While it is tempting to imagine that new techniques or methods will simply wipe away traditional practices, this is rarely the case for scientific revolutions [10] let alone for a field with the complexities of

healthcare entwined with the economics of industry and regulatory concerns. While it is beyond the scope of this chapter, an examination of the potential gains in health and economic terms to be achieved from a unification of the field across these areas is motivation enough to hold this out as a goal.

What follows is a proposed research agenda which concentrates on a few areas (1) that provide common ground among researchers, industry professionals, and regulators, (2) in which technological advances are beginning to provide significant advances, and (3) in which research informaticists can provide major contributions and guidance.

---

### Research Agenda for Modern Pharmacovigilance

**A Note on Machine Learning** Over the last few years, as computing power has reached sufficient levels and research has matured, there has been an explosion in the application of machine learning techniques to many areas in healthcare and pharmaceutical research [11–13].

Such is the meteoric rise in the use of machine learning and algorithmic computation across healthcare and research that research topics 3, 4, and 5 in the sections below are largely concerned with the impact in these areas, whereas just a few years ago, they would be mentioned in passing.

It is no longer possible to approach a research agenda for pharmacovigilance without careful consideration of how these techniques and technologies are changing what is possible. But while their impact is considered here in light of their impact on the field, this chapter makes no attempt to evaluate specific techniques in machine learning or artificial intelligence, except as they apply to the specific research topics described below.

## Topic 1: The Operational Definition of an Adverse Event

The regulatory definition of an adverse event (AE)<sup>1</sup> is well established, with the term coming into common use in the 1930s and being refined in the 1960s and 1970s, at the same time that formal pharmacovigilance systems began to be established [6]. There has been a refinement of the term since then, but the general definition has remained fairly stable. For our purposes, what is important to note is that the definition of an AE was conceived at a time when the Internet, social media, big data, and the promise of large amounts of digital healthcare data were nascent or nonexistent. The most important effect this has had on the definition of an AE is to cast it in terms of a paper metaphor—we picture in our minds collecting AEs onto forms, and we think of the various elements of the form, the amount of information to be collected, and the location of what type of information should go together, all in terms of a piece of paper. The insidious use of this metaphor encourages a habitual mode of thought which, having been ossified in regulatory definitions, is hard to escape. And while the metaphor has been extended significantly, initially to cover copies and facsimiles and later to include the concept of electronic data stores, the impact of the Internet and the wholesale digitization of healthcare data have stretched the paper metaphor to its limit. It is past time for a re-examination of the fundamental definitions of the field.

The need to update our concepts in regard to how we define AEs becomes evident when we seek to operationalize the definition of an AE in order to implement it into systems and use it for research. The classic operational definition derived originally from regulatory use is that a

valid adverse event report has “four elements”: an identifiable patient, an identifiable reporter, a suspect drug, and a serious adverse event or fatal outcome [14]. Over time the requirements for a regulatory report (which were created to help busy doctors understand what to report on a piece of paper) have become conflated with the definition of an AE, to the point where we might define a report that is missing these elements as irrelevant. But when we understand that the “four elements” are simply an operational definition meant to assist doctors in reporting, we can see that, given the digitization of healthcare data today, there is a need for a new operational definition.

An example illustrates the difficulties that arise from the mismatch of our concepts and the digital reality today in healthcare. In 2010, a pilot study demonstrated for the first time that it was possible to collect AEs at the point of care directly from an electronic health record, with minimal impact on clinicians, and to have those events sent electronically to FDA, in a matter of minutes after the initial recognition of the event [15]. At the time this study was performed, one of the authors engaged in fierce debate with industry colleagues over the fact that the individual physician’s name was masked on the report (although the medical institution was known) and therefore the report was not a “qualified” AE (personal communication). This arcane argument took place as a result of an outdated operational definition for an AE, so that even though we could infer the existence of an individual physician given the design and operation of the electronic health record, the exact requirement of an “identifiable reporter” could be interpreted to mean the report was disqualified.

Healthcare research has no such operational definition for what constitutes an AE, and while this allows for a more rational approach to collecting medically relevant information, it means that there can be no direct sharing of approaches or interpretation of findings between the different sectors. And the reason such operational

---

<sup>1</sup>Those familiar with the use of the term “ADR” (adverse drug reaction) vs “AE” (adverse event) should note that this discussion does not attempt to differentiate between those stricter definitions. Here the term “AE” is meant to be used in a general sense of a reported or noticed problem or concern.

definitions are required by regulators and industry is that there are massive efforts which span companies and continents, which require some semblance of uniformity if the attempts to perform pharmacovigilance are to yield useful results.

Given that both sectors have an interest in AEs, it would be of great benefit if a more inclusive, subtle, and encompassing operational definition of an AE could be developed. Informaticians seeking to make progress here could begin with sound medical concepts to define the broadest category of adverse events. Clearly this work should be built on existing useful clinical models and ontologies (a topic discussed later), but an understanding of the regulatory definitions will be important as well. The goal would be to create a continuum of definitions based on informatics rather than the incongruous set of definitions that exist today. In this way we can imagine that AEs of “regulatory interest” would be a subset of a larger group of medical interest.

It could be argued that this distinction exists today—AEs collected as a matter of course in healthcare are examined to see if they meet regulatory criteria, and if so, they are classed as such. The problem with this approach is that using the outdated “four elements” to define AEs of regulatory interest ignores a significant number of medically interesting events. The time has come to rework the operational definition to better align with what qualifies today as an AE from work being done by researchers in healthcare.

This topic takes on greater urgency today as the use of “real-world data” becomes more commonplace in regulated clinical research which opens up the consideration of the high-dimension, longitudinal data found in electronic health records and sourced from wearable sensors. The richness of the available data demands a more nuanced and expanded definition for adverse events.

## Topic 2: Expanding and Formalizing the Data Model

Similar to the operational definition of an AE, the data model used to report AEs was developed

from a need by regulators to have industry be able to report, in a consistent manner, AE reports. The original document of the 1996 document from the International Council of Harmonization (ICH) that addressed the “Data Elements for Transmission of Individual Case Safety Reports” (ICSR) was designated “E2” (the ICH designation for pharmacovigilance documents) and “B” referred to the particular document that defined data elements [16]. Hence, when referring to “E2B,” we are referring to the underlying data model for an AE.

The E2B data model is well-developed and used internationally, which is an advantage. But as is the case with the operational definition of an AE, E2B had its origins long before big data, the Internet, and the dramatic increase in digitized healthcare data. With the most recent version (E2BR3), the overall standard is based upon a HL7 ICSR model that is capable of supporting the exchange of messages for a wide range of product types (e.g., human medicinal products, veterinary products, medical devices). This is an excellent move toward more functionality within the regulatory reporting realm, but whereas this works well to allow submission of AEs to regulators, from an informatics perspective, looking to the future support of research across healthcare, this is lacking.

Contrast this with the type of large-scale research done today using very large and disparate datasets. This work has driven the creation of common data models which often include adverse events. A good example of this is the Observational Medical Outcomes Pilot (OMOP) common data model (CDM) [17] produced by OHDSI (Observational Health Data Sciences and Informatics). The OMOP CDM was created to use in the systematic analysis of disparate observational databases, and to this end it has a common format and common terminologies, vocabularies, and coding schemes.

Use of this approach in pharmacovigilance is what Koutkias and Jaulent have called the “computational approach” [18], in this case specifically for signal detection. The authors argue that pharmacovigilance should exploit all possible sources of information that may impact drug and

device safety, and they do an excellent job of reviewing the sources, tools, and approaches. Most importantly, they suggest that semantic technologies are the right approach to this new pursuit of using diverse data sources in a unified fashion.

One semantic technology increasingly popular in clinical informatics is ontologies—explicit, formal specifications of terms or concepts in a domain and the relationships among them [19]. An early introduction of ontologies to the field of pharmacovigilance came in 2006 when Henegar et al. looked at formalizing MedDRA, the standardized medical terminology used for international regulatory purposes, one of which is to report AEs [20]. What Henegar discovered with MedDRA is illustrative of many models and terminologies in use with pharmacovigilance—there were no formal definitions of terms in MedDRA, and this meant that no formal description logic could be applied to reason against data described with this terminology. The lack of formal logic and rigorous concept representation meant that inference was not possible based on semantic content.

For many years, those engaged in pharmacovigilance research in industry were well aware of the lack of a semantic layer, but it was considered simply an artifact of the way in which data was collected. Groupings and counts of terms in MedDRA were gathered, and what then followed was a long and arduous process of in effect manually applying the semantic layer back to the data. Ontologies have been demonstrated to significantly improve this situation and allow us to imagine the ability to combine large and disparate sources of data and properly infer from them [18, 20–22].

The challenge today is that there is still relatively sparse communication between the regulatory-facing tools used in pharmacovigilance and those being borrowed from computational biology and other disciplines allowing us to expand the data sources and techniques used in researching the safety of medical products. The SALUS study [23] took on the challenge of harmonizing data models and terminologies in an effort not typical in signal verification studies.

This approach holds great promise and engenders a significant amount of research, but SALUS was unusual in that the authors sought to harmonize the work with regulatory requirements. To achieve this, in addition to creating a rich ontology to work with the EHR, they mapped certain elements onto the previously described reporting standard, E2B (R2). And while this was an effective demonstration that it is possible to unify the healthcare, industry, and regulatory needs in pharmacovigilance (by seeking a logical lower-level ontological representation), the fact that now a major revision to E2B (R3) has come into effect and demonstrates the continued balkanized nature of the field.

There is no lack of definitions for adverse events. The ongoing development of the FHIR standard [24] has generated renewed interest in this area, as well as ongoing work on the OMOP standard [25].

Work by informaticists is needed to unify and maintain the representations needed in pharmacovigilance, and settling on a set of key ontologies would be a dramatic step forward and would enable better utilization of diverse sources of data, more economical translation of data for industrial research, and more accurate, better-quality communication of this information for regulatory purposes. The field of oncology research may be a useful model for informaticists looking to improve the definitions in pharmacovigilance. As a result of the dramatic increase in genomic data and other real-world data, oncology has been learning to manage massive amounts of detailed data with precision medical concepts—and so is often at the forefront of informatics work that impacts clinical research. Becoming familiar with the unique approach to toxicities, adverse event definitions, and attempts to reconcile healthcare and research concepts in oncology is an excellent introduction to possible solutions [26, 27].

### Topic 3: Terminologies

Since the beginning of medical and industrial research, terminologies have been developed in an attempt to categorize and standardize work.

And it has long been recognized that the problem of semantics, or the meaning of terms in medicine and healthcare research, cannot be fully divorced from the terminologies used to describe things [28, 29]. Along with heterogeneous data models, lack of consistency in various terminologies and how they're applied has been a challenge even before described succinctly by Cimino and is understood as a lynchpin to using EHRs for big data research [30].

Recently, the work being done in machine learning, ontologies, and computational methods is shedding new light on ways to tame the terminology issues, such that it is now imaginable that the problem of inconsistency could be solved by a logically rigorous ontology which binds terminologies to data models [31]. As a discussion of ontologies preceded this section (see Chap. 19), here we highlight work being done in machine learning which impacts challenges with terminologies.

For the last several years, researchers have looked at computer-assisted ways to extract AEs from text (specifically from narratives in AE reports) [32], but more recently new levels of sophistication in handling terminology as part of the process have been demonstrated. Jiang et al. evaluated using machine learning-based approaches to extract clinical entities from hospital discharge summaries written in free text [33]. Clinical entities included medical problems, tests, and treatments. While this work did not specifically address identification of AEs, the clinical and conceptual challenges are the same, and indeed in some cases, medical problems are adverse events.

Of interest was their finding that traditional mapping of text to controlled vocabularies (time-consuming work that often reflects individual preference) could be helped by accurate boundary detection by machine learning systems which do named entity recognition (NER) tasks (find and classify words and phrases into semantic classes). They hypothesize this system could help recognize unknown words based on context and so could supplement traditional dictionary-based NLP systems. The implication here is that the task of finding and accurately coding adverse

events (among other medical concepts) could be significantly standardized and automated via the methods described.

For pharmacovigilance, this would have a direct application not only in finding AEs in discharge summaries but also in recognizing AEs from patient diaries and notes, where an expression that refers to an AE may have no recognition in a dictionary-based system (e.g., “this stuff split my head into”—where the vernacular refers to a drug-induced headache, but the terms and the misuse of “into” vs “in two” makes machine recognition challenging).

The development of a machine learning approach demands better-defined, more logically consistent datasets, and this has spurred work which will change the traditional challenges associated with terminologies. Borrowing from a bioinformatics and systems biology approach, Cai et al. created ADReCS—the Adverse Drug Reaction Classification System [34]. ADReCS is an ontology of AE terms built with MedDRA and UMLS with hierarchical classification and digital identifiers. This means that direct computation on ADR terms can be achieved using the system, a significant step for the efficient use of machine learning technologies. We can imagine a future where this system or ones similar are expanded and mapped to other ontologies built in a similar manner, allowing for an approach to pharmacovigilance that is unlike anything in the past. As we reach this stage of computational maturity in pharmacovigilance, it will create a very significant driver for the biopharmaceutical industry, which spends a great deal on gathering data from disparate sources to test drug safety hypotheses and to standardize and recode that data into common formats that can be submitted to regulators. As systems like ADReCS become the norm, many of the inefficiencies the industry now faces will begin to disappear.

As with ontologies, work is needed to expand the most promising systems and to find the most universal and effective representations of terminologies that can migrate successfully from healthcare to industry to regulators with no loss of meaning and will decreased manual effort.



#### Topic 4: Discovery/Curation of AEs

Research on the discovery of AEs is being done in every possible source—electronic health records, social media, registries, large databases, real-world data from insurance claims, and other sources [35]. In 2012, Harpaz et al. set the stage for the use of novel methodologies using large datasets with their review of current work [36]. The authors made several salient points regarding the new research methods, including the fact that (1) combining data from heterogeneous sources requires the development of new and reproducible methods, (2) standardized (and simulated) datasets will grow in importance to allow rapid testing of new methods, and (3) standards in PV must be developed to evaluate algorithmic approaches applied to the data. In 2013 Jiang et al. began work on ADEpedia 2.0, which built on their previous AE knowledge base derived from drug product labels; in keeping with the direction laid out by Harpaz, in 2.0 the authors began to enrich the database with data from UMLS (Unified Medical Language System) and EHR data, with a goal to create a standardized source of AE knowledge [37]. Banda et al. continued this approach, standardizing the FDA's FAERS (FDA Adverse Event Reporting System) database [38]. They provided a curated database removing duplicate records, mapping the data to standardized vocabularies with drug names mapped to RxNorm concepts and outcomes mapped to SNOMED-CT concepts, and created a set of summary statistics about drug-outcome relationships for general consumption. While not involved directly with machine learning, this approach pointed the way toward further machine-based approaches by providing all source code for the work, so that it could be used and updated as needed, and by mapping outcomes and indications to SNOMED-CT, this allows for direct linkage to other ontologies.

Since that time, an explosion of work has taken place in all three areas identified by Harpaz, emphasizing the discovery of AEs using machine learning combined with statistical techniques [39–44].

The study by Bean et al. [39] serves to illustrate a new way of approaching discovery of AEs in the postmarketing phase—one that doesn't wait for a series of reports to emerge; rather it takes advantage of what until recently were infrequently connected sources of data to discover previously unknown AEs due to specific drugs and to validate this via EHRs. The authors constructed a knowledge graph with four primary sources of data: drugs, protein targets, indications, and adverse reactions that predicted AEs from public data. They then used this to develop a machine learning algorithm and deployed that algorithm on an EHR. The algorithm was fed by an NLP pipeline developed to parse free text in the EHR. This work is similar to work on prediction of AEs using structure-activity relationships [45], gene expression [46], and protein drug targets [47]. In this work we can see a computational biological approach which can view with the current biology-based approach that has paid dividends but dominated PV for decades.

In 2017 Voss et al. moved the field forward significantly with their work to automatically aggregate disparate sources of data into a single repository [48] that allows a machine learning approach to selecting positive and negative controls for pharmacovigilance research design testing. As previous work demonstrated, creating a reference database for pharmacovigilance using manual or even semi-manual methods, is extremely time- and resource-intensive. The authors built on previous work (described in Banda) and added the relationship between a drug and a health outcome of interest (HOI). They performed a quantitative assessment of how well the evidence base could discriminate between known positive drug-condition causal relationships and drugs known to be not associated with a condition, thus allowing the automated creation of an assessment for pharmacovigilance research study designs that allows comparisons across designs with a significant savings in time and increase in standardization. The authors worked through methods for accepting data from various sources at various granular levels, for example, mapping the source at either an ingredient level of a drug or at a

clinical drug level and subsequently aggregating evidence to individual ingredients to allow analysis across the dataset.

While work in AE discovery occurs at every level and is often the primary topic in other researched covered under other research topics such as terminologies, the topic of curation is not one typically addressed except in individual efforts that are not reproducible and rarely maintained due to the intense effort required. This shift toward automated curation across various data sources will prove to be an important stimulus to the computational approach in pharmacovigilance, allowing a much more rapid and standardized testing of research designs. In the near future, we can expect more reference sets which can be used to train machine learning algorithms and test large-scale analysis methods.

Just as the creation of high-quality curated datasets in machine learning is driving forward progress across many fields [49], we can expect the same to occur in pharmacovigilance as work continues. It is insightful to review the dramatic effect that a massive, well-curated (automatically generated) dataset can have on accuracy of machine learning algorithms in the example of ImageNet [50, 51], a database of over 14 million image URLs that are labeled to provide a curated set. Prior to establishing this dataset, progress in visual object recognition was steady but slow. In 2012, using a deep convolutional neural network trained on ImageNet, researchers bested other networks by over 40% to the next best [52]. This massive, curated dataset is widely attributed as one of the primary drivers of the deep learning revolution. Other such databases (VigiBase with 21 million Case safety reports in 2020) are emerging as resources to advance AE detection and even scientific discovery [53].

The moral of this story for pharmacovigilance is that a focus on the creation of large, curated, automatically created test datasets has the potential to move a computational approach to pharmacovigilance forward just as quickly if not more quickly than the best analytical methods. This is certainly an area for future informatics research.

## Topic 5: Delayed Toxicity and Complex Causal Assessments

The tragic discovery of delayed hepatotoxicity caused by fialuridine is required reading in any pharmacovigilance or clinical research education. It is important to understand just how difficult it was at that time to attribute observed toxicities to the drug, given how they initially presented in patients and the presence of similar symptoms due to underlying disease or caused by an initial therapeutic response. These challenges, coupled with the piecemeal accumulation of information over a period of time, made it difficult to form a conviction that fialuridine caused a fatal toxicity—although, as some argue, evidence was clearly present [54, 55]. The 1995 Institute of Medicine report on the review of events leading up to the tragic deaths of five patients in a 1993 clinical trial of fialuridine for hepatitis B concluded that overall, clinical researchers involved in various trials acted correctly and made the best decisions possible given the available information. Looking at the set of trials that were done over a period of several years, however, one cannot help but be struck by the series of “clues” pointing to fialuridine and how, when taken together, they provide a strong signal that the drug was implicated [56].

In our current post-behavioral economics atmosphere, it may be easier for us to appreciate how we could fail to recognize a problem of delayed toxicity in a drug: humans are superb at pattern recognition over a relatively short timeframe, but our skill degrades rapidly as cause is separated in time from effect and obscured by other possible causes. In pharmacovigilance, one has a feeling of inadequacy when it comes to sorting out the possible links between drugs and toxicities, except in the most obvious and common cases. The investigations into fialuridine-delayed toxicity produced better regulation and reasonable research recommendations [56], but beyond these improvements, not much has been gained in our ability to recognize delayed toxicity in drugs from complex situations.

A less dramatic but conceptually similar challenge faces anyone seeking to sort out what drugs may be contributing to a patient's clinical signs and symptoms when they have underlying disease and are on a multiple drug regimen. The classic questions regarding “dechallenge/rechallenge” (whether a sign or symptom stopped once drug was stopped, and returned after drug was restarted) and the time course of drug dose vs appearance of symptoms are well designed but often unanswerable in a real-world situation. Oncology trials come to mind as a particularly challenging environment in which to attribute cause to individual drugs.

These scenarios are not unique to pharmacovigilance. They share the same basic external challenges—incomplete information, competing causes, extended overtime, and internal challenges—idiosyncratic human perception, and bias with pursuits as diverse as cognitive psychology and behavioral economics [57] or the study of policy impacts [58].

Computational approaches to these questions hold out promise to provide the most significant advancement in years for pharmacovigilance, by transferring the burden of recognition to computers working with large datasets using sound methods. Most of the work reviewed earlier in the recognition of AEs applies here as well. Huang et al. systems pharmacology approach of combining clinical observation with molecular biology [42] can be seen as template for research in predicting toxicities in drugs and arming researchers with information that will enhance the design as well as the monitoring of trials using drugs with increasingly complex mechanisms of action. Recent similar work indicates that a systems pharmacology or computable biology approach holds out great promise in predicting toxicities at an earlier stage than previously imagined [59–62].

Combining data across disciplines in a computable framework is a fertile area of research, especially as it applies to predicting toxicities in a real-world setting. The contribution of informatics to this work can have a tangible and concrete impact in improving safety for patients.

Arming clinical researchers and pharmacovigilance professionals with these methods holds out hope that another fialuridine tragedy would be avoided today.

## Topic 6: Risk Profiling of the Individual

The concept of precision medicine that medical care can be tailored—especially in a genomic and molecular sense—to select groups of patients is now commonplace and being realized in the design of clinical trials and healthcare policy in addition to medical practice. In pharmacovigilance, however, there is a need for better understanding of how the concepts of precision medicine can be incorporated into goals and practice. This section simply poses some basic open questions that informaticians can help to address in order to improve the theoretical basis of pharmacovigilance. But while the ideas here are to some degree speculative, the authors believe they should be taken seriously, as they are at the heart of pharmacovigilance itself.

A simple coined term “precision pharmacovigilance” is enough to raise questions and spark ideas about how the discoveries in medicine and biology can be more directly taken up in the study of drug and device safety [63]. But a broader (and more provocative) research question to ask is *Is it possible to provide to an individual a ‘risk profile’ as it relates to their particular drug and/or device regimen?* One aspect of the question relates to the degree to which we can simply follow the discoveries in precision medicine and practice pharmacovigilance along the way—e.g., looking at AEs in certain genetic subgroups while undergoing treatment with immune modulators. It could be said that in this respect, there's nothing new here; pharmacovigilance has looked at subgroups of patients for some time [64] and continues to bear fruit [65].

But recent research in methods dealing with large-scale longitudinal observational databases [66, 67] allows us to imagine a scenario different

from that of looking at the AEs related to certain subtypes of patients—what if we could predict the risk of being a certain person (age, race, genetic makeup), taking a certain set of drugs (let’s say a regimen of five separate drugs), living in a certain area of the world, and having a particular occupation? Can we reach the point where we can tell you that for you as an individual, you have a 60% chance of a significant toxicity if you fit the above profile? The question serves less to examine how much data would it take to provide an exact answer and more to challenge us to decide how feasible it is to pursue this goal. Can pharmacovigilance aspire to studying and predicting risk not only for patient subtypes but for situational circumstances?

At this early stage of discovery and application in big data, machine learning, and improving methods, it is important to keep an open mind about what pharmacovigilance can become. Being able to speak directly to select groups of patients who are living in specific circumstances as regards their drug therapy was an original motivation for pharmacovigilance and we believe should continue to inspire research.

### **Topic 7: Emerging Data and Technologies for Pharmacovigilance**

Since the first edition of this chapter, the most significant change to impact the clinical research informatics of pharmacovigilance is the continuation of the trends in all of the related research in AI and the continued increase in the amount and availability of data. Such is the pace of change that the authors expect any future discussion of pharmacovigilance will need to include these topics as part of the primary discussion.

Several of the more important developments are highlighted here.

#### **Interoperability of Healthcare Data**

At the time of this writing, there is a major shift toward healthcare data becoming “interoperable” [68]. This is the long-sought-after goal of making healthcare data available and portable to improve

healthcare delivery, patient care, and research. The confluence of more modern, web-based standards such as FHIR, the regulatory push from the Office of the National Coordinator implementing certain requirements of the 21st Century Cures Act, most specifically the “Interoperability Rule” [69] has given rise to a new breed of “interoperability vendors” who are connecting data across healthcare systems and in the process making more standardized healthcare data available for clinical and regulated research. This trend promises to help solve one of the more intractable problems in scaling clinical research using real-world data—the heterogeneous technical and data environment that exists across the United States and globally.

#### **Alphafold 2**

Any researchers interested in novel methods to improve pharmacovigilance should gain a deep understanding of the seminal and dramatic achievement of Alphafold 2—the deep learning AI system developed by Alphabet’s/Google’s DeepMind which performs predictions of protein structure and which basically solved the protein folding problem that was a major goal of researchers and drug developers [70–72].

The obvious impact of Alphafold 2 for the drug development industry is clear as it dramatically reduces the time needed to explore structure-function relationships of drug molecules or their molecular targets. This has sparked a dramatic uptake in the pursuit of practical applications across the industrial and academic worlds.

But the larger implications of Alphafold 2 for pharmacovigilance are apparent when we examine how the approach to specific knowledge domains coupled with recent deep learning designs can produce results that were previously thought impossible. This will have direct application not only to the molecular basis of adverse events (especially where molecular biomarkers are associated) but also portends a new, data-focused approach to research in pharmacovigilance that is entirely foreign to the past event-based approaches. This is a nascent area of research to consider for pharmacovigilance specifically, but the authors suspect that direct

application of the system designs used for AlphaFold 2 will become the next horizon of research.

### **Transformer-Based Language Models and GPT3**

If AlphaFold 2 applicability to pharmacovigilance seems to be in the misty distance, Generative Pre-trained Transformer 3 (GPT3) and its impact in just the last few years provide a clear demonstration of the direction of new research. Since their introduction, transformer-based large language models have shown an unusual and not predicted capability to solve problems across various domains [73]. Clearly the most dramatic discovery was that such “language” models can actually solve task-agnostic problems, meaning that they can be used to learn images and associated text with pictures, actions, and more [74].

Research is underway to apply language models to the challenges of pharmacovigilance. Guan and Devarakonda used BERT (Bidirectional Encoder Representations from Transformers) to find adverse events and the drugs that caused them in the literature, improving on previous NLP approaches [75]. Wang et al. go further by combining transformer-based language models with Judea Pearl’s “do-calculus” method of causality assessment [76, 77]. As a result they begin to breach the age-old problem of causality in adverse event assessment [78].

As these examples illustrate, a continuation of the general AI, machine learning, natural language processing trends have reached the stage where their power and applicability as research designs cannot be denied. We expect this research to proliferate and begin to address fundamental issues in pharmacovigilance.

Empowered with tools and knowledge and strong patient advocacy networks, the role of the patient in identifying and reporting potential adverse events will inevitably grow—and with it continue to change the paradigm of AE detection, investigation, and action. This, of course, goes against the very ingrained mindset of regulators, who see population approaches as the only PV worthwhile, relegating personal issues to physicians. This current anachronistic attitude toward

PV is ironic, given the parallel trumpeting and funding of personalized medicine and research. As access and usability of technology continues to evolve and support increased patient engagement in both healthcare and research, attitudes and paradigms around PV will likely evolve, as will approaches and regulations to understand and improve the safety of medications, vaccines, and medical devices.

---

## **The Future of Pharmacovigilance**

The future of pharmacovigilance lies, as with many fields, in the application of AI methods to increasingly large datasets. There are now two frontiers addressing this work in the larger context of drug safety. The first frontier has been described above, and it is applying new AI methods and models directly to fundamental challenges in pharmacovigilance. The second frontier represents the industry response to this undeniable trend. Taking a more conservative approach (as can be expected given the regulatory requirements of the field), there is now great interest in how moderate machine learning and NLP approaches can be applied to the administrative and labor-intensive processes in industrial pharmacovigilance work [79–81].

It is likely that machine learning and NLP will provide valuable time and cost savings to current industry processes, and this can only be positive for pharmacovigilance as a whole. It does not, however, address the more fundamental questions that AI has raised for industrial pharmacovigilance just as it has for most other industries. Ball and Dal Pan provide the quintessential example of this approach in their excellent review article “Artificial Intelligence” for Pharmacovigilance: Ready for Prime Time?” [82]. They review the current regulated process for pharmacovigilance and discuss points at which AI (more specifically a truncated definition of “algorithms”) could be useful and where there are still challenges. The very approach to “algorithms” in the article belies the fact that the authors limit the discussion by employing an outdated definition of what AI is based on current

work as described above. To call GP-3 and “algorithm” loses all meaning when you consider the full version has over 170 billion learning parameters. Just as the initial digitization of adverse event data suffered from an outdated paper metaphor, which limited its usefulness and slowed meaningful change, so the metaphor of an “algorithm” representing AI is causing the same issues in the professional field of pharmacovigilance. We hope that the half-life of this misrepresentation will be much shorter than was the paper metaphor for data. Finding accurate definitions, ontologies, terminologies, and approaches to pharmacovigilance based on the methods and proven results being discovered today in domain-specific data sets combined with large language models is work in tremendous need for the contributions of clinical research informaticists.

## Conclusion

The above models of observational research and signal detection in the real world will necessitate standard data representations—including controlled terminology and shared data, information, and (formal) knowledge models. These challenges are nontrivial and are common obstacles for other major informatics activities, such as improving data exchange, collection of longitudinal data, and real-time clinical decision support, which is the holy grail of informatics and electronic health. Because standard data models and terminology are central to so many EHR goals and stakeholders, lots of energy and resources are directed here, and hence good reason to be optimistic that standardized data from EHR systems will one day be available to support pharmacovigilance. These same stakeholders can/will also make the case (to the public and healthcare consumer) of the vital role that standardized and quality clinical data will play in public health. These stakeholders (and drivers, from industry, patient advocates, and the public) can highlight pharmacovigilance as a *public health issue*—and one that is relevant to clinical care as well as research.

In the end, the divergence of patient-facing drug safety and industrial pharmacovigilance continues today as it was first described in the original publication of this chapter. While the regulated industry provides a certain stability, we must look to innovations in healthcare approaches to risk assessment and personalized medicine for our future models of pharmacovigilance to bring significant improvements to patients’ lives and safety. As Sir Tim Berners-Lee, who invented the World Wide Web in 1989, noted:

*Data is a precious thing and will last longer than the systems themselves.*

## References

1. Härmark L, van Grootheest AC. Pharmacovigilance: methods, recent developments and future perspectives. *Eur J Clin Pharmacol.* 2008;64(8):743–52. <https://doi.org/10.1007/s00228-008-0475-9>. Epub 2008 Jun 4. <https://www.ncbi.nlm.nih.gov/pubmed/18523760>.
2. van Grootheest AC, Richesson RL. Pharmacovigilance. In: Richesson R, Andrews J, editors. *Clinical research informatics. Health informatics.* London: Springer; 2012. [https://doi.org/10.1007/978-1-84882-448-5\\_19](https://doi.org/10.1007/978-1-84882-448-5_19).
3. Friedman LM, Furberg CD, DeMets DL. *Assessing and reporting adverse events.* In: *Fundamentals of clinical trials.* New York: Springer; 1998. p. 170–84.
4. Andrews EB, Moore N. *Mann’s pharmacovigilance.* 3rd ed. Chichester: Wiley-Blackwell; 2014.
5. World Health Organization. What is pharmacovigilance? [http://www.who.int/medicines/areas/quality\\_safety/safety\\_efficiency/pharmvigi/en/](http://www.who.int/medicines/areas/quality_safety/safety_efficiency/pharmvigi/en/). Accessed 4 Feb 2022.
6. Talbot J, Aronson JK, editors. *Stephens’ detection and evaluation of adverse drug reactions: principles and practice.* 6th ed. Chichester: Wiley; 2011.
7. Ibara MA, Richesson RL. Back to the future: the evolution of pharmacovigilance in the age of digital healthcare. In: Richesson R, Andrews J, editors. *Clinical research informatics. Health informatics.* Cham: Springer; 2019. [https://doi.org/10.1007/978-3-319-98779-8\\_20](https://doi.org/10.1007/978-3-319-98779-8_20).
8. Coase RH. The nature of the firm. *Economica.* 1937;4(16):386. <https://onlinelibrary.wiley.com/doi/full/10.1111/j.1468-0335.1937.tb00002.x>. Accessed 4 Feb 2022.
9. Naughton J. How a 1930s theory explains the economics of the internet. *The Guardian.* 7 Sept 2013. <http://www.theguardian.com/technology/2013/sep/08/1930s-theory-explains-economics-internet>. Accessed 4 Feb 2022.

10. Kuhn TS. The structure of scientific revolutions. Chicago: University of Chicago Press; 2012. <http://www.press.uchicago.edu/ucp/books/book/chicago/S/bo13179781.html>.
11. Chilcott M. How data analytics and artificial intelligence are changing the pharmaceutical industry. *Forbes Magazine*. 10 May 2018. <https://www.forbes.com/sites/forbestechcouncil/2018/05/10/how-data-analytics-and-artificial-intelligence-are-changing-the-pharmaceutical-industry/>.
12. Dua S, Rajendra Acharya U, Dua P. Machine learning in healthcare informatics. *Intelligent systems reference library*. 2014. <https://link.springer.com/book/10.1007%2F978-3-642-40017-9>.
13. Jiang F, Jiang Y, Zhi H, Dong Y, Li H, Ma S, Wang Y, Dong Q, Shen H, Wang Y. Artificial intelligence in healthcare: past, present and future. *Stroke Vasc Neurol*. 2017;2:230–43. <https://doi.org/10.1136/svn-2017-000101>.
14. U.S. Food and Drug Administration. Guidance for industry postmarketing adverse event reporting for nonprescription human drug products marketed without an approved application. Oct 2007. <https://www.fda.gov/downloads/drugs/guidancecomplianceregulatoryinformation/guidances/ucm071982.pdf>.
15. Linder JA, Haas JS, Iyer A, Labuzetta MA, Ibara M, Celeste M, Getty G, Bates DW. Secondary use of electronic health record data: spontaneous triggered adverse drug event reporting. *Pharmacoepidemiol Drug Saf*. 2010;19(12):1211–5. <https://doi.org/10.1002/pds.2027>.
16. Research. Center for Drug Evaluation and Guidances (Drugs). E2B(R3) electronic transmission of individual case safety reports implementation guide—data elements and message specification; and appendix to the implementation guide—backwards and forwards compatibility. n.d. <https://www.fda.gov/drugs/guidancecomplianceregulatoryinformation/guidances/ucm274966.htm>.
17. OMOP. Common data model—OHDSI. n.d. <https://www.ohdsi.org/data-standardization/the-common-data-model/>. Accessed 8 Mar 2018.
18. Koutkias VG, Jaulent M-C. Computational approaches for pharmacovigilance signal detection: toward integrated and semantically-enriched frameworks. *Drug Saf*. 2015;38(3):219–32.
19. Gruber TR. A translation approach to portable ontology specifications. *Knowl Acquis*. 1993;5(2):199–220.
20. Henegar C, Bousquet C, Louët AL-L, Degoulet P, Jaulent M-C. Building an ontology of adverse drug reactions for automated signal generation in pharmacovigilance. *Comput Biol Med*. 2006;36(7):748–67.
21. Pacaci A, Gonul S, Anil Sinaci A, Yuksel M, Erturkmen GBL. A semantic transformation methodology for the secondary use of observational healthcare data in postmarketing safety studies. *Front Pharmacol*. 2018;9:435.
22. Personeni G, Bresso E, Devignes M-D, Dumontier M, Smail-Tabbone M, Coulet A. Discovering associations between adverse drug events using pattern structures and ontologies. *J Biomed Semant*. 2017;8(1):29.
23. Yuksel M, Gonul S, Erturkmen GBL, Sinaci AA, Invernizzi P, Facchinetti S, Migliavacca A, Bergvall T, Depraetere K, De Roo J. An interoperability platform enabling reuse of electronic health records for signal verification studies. *Biomed Res Int*. 2016;2016:1–18. <https://doi.org/10.1155/2016/6741418>.
24. Overview—FHIR v4.3.0. n.d. <https://www.hl7.org/fhir/overview.html>. Accessed 11 Sept 2022.
25. OMOP common data model. n.d. <https://www.ohdsi.org/data-standardization/the-common-data-model/>. Accessed 11 Sept 2022.
26. Le-Rademacher JG, Hillman S, Storrack E, Mahoney MR, Thall PF, Jatoi A, Mandrekar SJ. Adverse event burden score—A versatile summary measure for cancer clinical trials. *Cancers*. 2020;12(11) <https://doi.org/10.3390/cancers12113251>.
27. Ewer MS, Herson J. Cardiovascular adverse events in oncology trials: understanding and appreciating the differences between clinical trial data and real-world reports. *Cardiooncology (London, England)*. 2022;8(1):13.
28. Cimino JJ, Clayton PD, Hripcsak G, Johnson SB. Knowledge-based approaches to the maintenance of a large controlled medical terminology. *J Am Med Inform Assoc*. 1994;1(1):35–50.
29. Schroll JB, Maund E, Götzsche PC. Challenges in coding adverse events in clinical trials: a systematic review. *PLoS One*. 2012;7(7):e41174.
30. Reich C, Ryan PB, Stang PE, Rocca M. Evaluation of alternative standardized terminologies for medical conditions within a network of observational healthcare databases. *J Biomed Inform*. 2012;45(4):689–96.
31. Ethier J-F, Dameron O, Curcin V, McGilchrist MM, Verheij RA, Arvanitis TN, Taweel A, Delaney BC, Burgun A. A unified structural/terminological interoperability framework based on LexEVS: application to TRANSFoRm. *J Am Med Inform Assoc*. 2013;20(5):986–94.
32. Kovacevic A, Dehghan A, Filannino M, Keane JA, Nenadic G. Combining rules and machine learning for extraction of temporal expressions and events from clinical narratives. *J Am Med Inform Assoc*. 2013;20(5):859–66.
33. Jiang M, Chen Y, Mei L, Trent Rosenbloom S, Mani S, Denny JC, Hua X. A study of machine-learning-based approaches to extract clinical entities and their assertions from discharge summaries. *J Am Med Inform Assoc*. 2011;18(5):601–6.
34. Cai M-C, Xu Q, Pan Y-J, Pan W, Ji N, Li Y-B, Jin H-J, Liu K, Ji Z-L. ADReCS: an ontology database for aiding standardization and hierarchical classification of adverse drug reaction terms. *Nucleic Acids Res*. 2015;43(D1):D907–13.
35. Murff HJ, Patel VL, Hripcsak G, Bates DW. Detecting adverse events for patient safety research: a review of current methodologies. *J Biomed Inform*. 2003;36(1–2):131–43.

36. Harpaz R, DuMouchel W, Shah NH, Madigan D, Ryan P, Friedman C. Novel data-mining methodologies for adverse drug event discovery and analysis. *Clin Pharmacol Ther.* 2012;91(6):1010–21.
37. Jiang G, Liu H, Solbrig HR, Chute CG. ADEpedia 2.0: integration of normalized adverse drug events (ADEs) knowledge from the UMLS. *AMIA Joint Summits Transl Sci Proc.* 2013;2013:100–4.
38. Banda JM, Lee E, Vanguri RS, Tatonetti NP, Ryan PB, Shah NH. A curated and standardized adverse drug event resource to accelerate drug safety research. *Sci Data.* 2016;3:160026.
39. Bean DM, Honghan W, Iqbal E, Dzahini O, Ibrahim ZM, Broadbent M, Stewart R, Dobson RJB. Knowledge graph prediction of unknown adverse drug reactions and validation in electronic health records. *Sci Rep.* 2017;7(1):16416.
40. Boland MR, Jacunski A, Lorberbaum T, Romano JD, Moskovitch R, Tatonetti NP. Systems biology approaches for identifying adverse drug reactions and elucidating their underlying biological mechanisms. *Wiley Interdiscip Rev Syst Biol Med.* 2016;8(2):104–22.
41. Ho T-B, Le L, Thai DT, Taewijit S. Data-driven approach to detect and predict adverse drug reactions. *Curr Pharm Des.* 2016;22(23):3498–526.
42. Huang L-C, Wu X, Chen JY. Predicting adverse side effects of drugs. *BMC Genomics.* 2011;12(5):S11.
43. Jamal S, Goyal S, Shanker A, Grover A. Predicting neurological adverse drug reactions based on biological, chemical and phenotypic properties of drugs using machine learning models. *Sci Rep.* 2017;7(1):872.
44. Zhang W, Liu F, Luo L, Zhang J. Predicting drug side effects by multi-label learning and ensemble learning. *BMC Bioinform.* 2015;16:365.
45. Frid AA, Matthews EJ. Prediction of drug-related cardiac adverse effects in humans—B: use of QSAR programs for early detection of drug-induced cardiac toxicities. *Regul Toxicol Pharmacol.* 2010;56(3):276–89.
46. Wang Z, Clark NR, Ma'ayan A. Drug-induced adverse events prediction with the LINCS L1000 data. *Bioinformatics.* 2016;32(15):2338–45.
47. Pérez-Nuño VI, Souchet M, Karaboga AS, Ritchie DW. GESSE: predicting drug side effects from drug–target relationships. *J Chem Inf Model.* 2015;55(9):1804–23.
48. Voss EA, Boyce RD, Ryan PB, van der Lei J, Rijnbeek PR, Schuemie MJ. Accuracy of an automated knowledge base for identifying drug adverse reactions. *J Biomed Inform.* 2017;66:72–81.
49. Wikipedia contributors. List of datasets for machine learning research. Wikipedia, the free encyclopedia. 1 Jul 2018. [https://en.wikipedia.org/w/index.php?title=List\\_of\\_datasets\\_for\\_machine\\_learning\\_research&oldid=848338519](https://en.wikipedia.org/w/index.php?title=List_of_datasets_for_machine_learning_research&oldid=848338519).
50. ImageNet Large Scale Visual Recognition Competition (ILSVRC). n.d. <http://www.image-net.org/challenges/LSVRC/>. Accessed 2 Jul 2018.
51. Wikipedia contributors. ImageNet. Wikipedia, the free encyclopedia. 21 Jun 2018. <https://en.wikipedia.org/w/index.php?title=ImageNet&oldid=846928201>.
52. Gershgor D. The data that transformed AI research—and possibly the world. *Quartz.* 26 Jul 2017. <https://qz.com/1034972/the-data-that-changed-the-direction-of-ai-research-and-possibly-the-world/>.
53. Bihan K, Bénédicte LV, Funck-Brentano C, Salem J. Uses of pharmacovigilance databases: an overview. *Therapie.* 2020;75(6):591–8. <https://doi.org/10.1016/j.therap.2020.02.022>.
54. Bari A. Severe toxicity of fialuridine (FIAU). *N Engl J Med.* 1996;334(17):1135; author reply 1137–8.
55. The cure that killed. *DiscoverMagazine.com*. Discover Magazine. n.d. <http://discovermagazine.com/1994/mar/thecurethathill345>. Accessed 4 Jul 2018.
56. Institute of Medicine (US). Committee to review the fialuridine (FIAU/FIAC) clinical trials. In: Manning FJ, Swartz M, editors. *Review of the fialuridine (FIAU) clinical trials*. Washington, DC: National Academies Press (US); 1995.
57. Stiensmeier-Pelster J, Heckhausen H. Causal attribution of behavior and achievement. In: Heckhausen J, Heckhausen H, editors. *Motivation and action*. Cham: Springer International Publishing; 2018. p. 623–78.
58. NoNIE guidance on impact evaluation. World Bank Group. n.d. [http://siteresources.worldbank.org/EXT/OED/Resources/nonie\\_guidance.pdf](http://siteresources.worldbank.org/EXT/OED/Resources/nonie_guidance.pdf).
59. Ai H, Chen W, Zhang L, Huang L, Yin Z, Hu H, Zhao Q, Zhao J, Liu H. Predicting drug-induced liver injury using ensemble learning methods and molecular fingerprints. *Toxicol Sci.* 2018;165:100. <https://doi.org/10.1093/toxsci/kfy121>.
60. Kim E, Nam H. Prediction models for drug-induced hepatotoxicity by using weighted molecular fingerprints. *BMC Bioinform.* 2017;18(7):227.
61. Kotsampasakou E, Montanari F, Ecker GF. Predicting drug-induced liver injury: the importance of data curation. *Toxicology.* 2017;389:139–45.
62. Yang H, Sun L, Li W, Liu G, Tang Y. In silico prediction of chemical toxicity for drug design using machine learning methods and structural alerts. *Front Chem.* 2018;6:30.
63. Edwards V. Looking forward. *Pharmacovigilance in the next 5 years. The Industry Vision*. Available at: [https://www.ema.europa.eu/en/documents/presentation/presentation-pharmacovigilance-next-5-years-industry-vision-v-edwards\\_en.pdf](https://www.ema.europa.eu/en/documents/presentation/presentation-pharmacovigilance-next-5-years-industry-vision-v-edwards_en.pdf). Last access date 23 Sept 2022.
64. Lynch T, Price A. The effect of cytochrome P450 metabolism on drug response, interactions, and adverse effects. *Am Fam Physician.* 2007;76(3):391–6.
65. St Sauver JL, Olson JE, Roger VL, Nicholson WT, Black JL 3rd, Takahashi PY, Caraballo PJ, et al. CYP2D6 phenotypes are associated with adverse outcomes related to opioid medications. *Pharmacog Personal Med.* 2017;10:217–27.



66. Moghaddass R. The factorized self-controlled case series method: an approach for estimating the effects of many drugs on many outcomes. n.d.
67. Shaddox TR, Ryan PB, Schuemie MJ, Madigan D, Suchard MA. Hierarchical models for multiple, rare outcomes using massive observational healthcare databases. *Stat Anal Data Min.* 2016;9(4):260–8.
68. The Office of the National Coordinator for Health Information Technology. Interoperability. n.d. <https://www.healthit.gov/topic/interoperability>. Accessed 11 Sept 2022.
69. *ONC\_Cures\_Act\_Final\_Rule\_03092020.pdf*. n.d. [https://www.healthit.gov/sites/default/files/page2/2020-03/ONC\\_Cures\\_Act\\_Final\\_Rule\\_03092020.pdf](https://www.healthit.gov/sites/default/files/page2/2020-03/ONC_Cures_Act_Final_Rule_03092020.pdf). Accessed 23 Sept 2022.
70. Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, Tunyasuvunakool K, et al. Highly accurate protein structure prediction with AlphaFold. *Nature.* 2021;596(7873):583–9.
71. Skolnick J, Gao M, Zhou H, Singh S. AlphaFold 2: why it works and its implications for understanding the relationships of protein sequence, structure, and function. *J Chem Inf Model.* 2021;61(10):4827–31.
72. Heaven WD. DeepMind’s protein-folding AI has solved a 50-year-old grand challenge of biology. *MIT Technology Review.* 30 Nov 2020. <https://www.technologyreview.com/2020/11/30/1012712/deepmind-protein-folding-ai-solved-biology-science-drugs-disease/>.
73. Marr B. What is GPT-3 and why is it revolutionizing artificial intelligence? *Forbes Magazine.* 5 Oct 2020. <https://www.forbes.com/sites/bernard-marr/2020/10/05/what-is-gpt-3-and-why-is-it-revolutionizing-artificial-intelligence/>.
74. Brown TB, Mann B, Ryder N, Subbiah M, Kaplan J, Dhariwal P, Neelakantan A, et al. Language models are few-shot learners. 2020. arXiv [cs.CL]. arXiv. <https://proceedings.neurips.cc/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf>.
75. Guan H, Devarakonda M. Leveraging contextual information in extracting long distance relations from clinical notes. n.d.
76. Pearl J. An introduction to causal inference. *Int J Biostatistics.* 2010;6(2):7. <https://doi.org/10.2202/1557-4679.1203>.
77. Pearl J, Mackenzie D. *The book of why: the new science of cause and effect.* Basic Books; 2018.
78. Wang X, Xiaowei X, Tong W, Roberts R, Liu Z. InferBERT: a transformer-based causal inference framework for enhancing pharmacovigilance. *Front Artif Intell.* 2021;4:659622.
79. Bate A, Hobbiger SF. Artificial intelligence, real-world automation and the safety of medicines. *Drug Saf.* 2021;44(2):125–32.
80. Aronson JK. Artificial intelligence in pharmacovigilance: an introduction to terms, concepts, applications, and limitations. *Drug Saf.* 2022;45(5):407–18.
81. Schmider J, Kumar K, LaForest C, Swankoski B, Naim K, Caubel PM. Innovation in pharmacovigilance: use of artificial intelligence in adverse event case processing. *Clin Pharmacol Ther.* 2019;105(4):954–61.
82. Ball R, Dal Pan G. ‘Artificial intelligence’ for pharmacovigilance: ready for prime time? *Drug Saf.* 2022;45(5):429–38.