Artificial Intelligence (AI) in Pharmacovigilance: do we really need it…?

Part 1: The Technicalities

Initial version July 2019 - Reviewed and Updated on 21 September 2021

Conflict of Interest and Authorship Conformation Form

- All authors have participated in drafting the article or revising it critically for important intellectual content.
- This manuscript has not been submitted to, nor is under review at, another journal or other publishing venue.
- The authors have no affiliation with any organization with a direct or indirect financial interest in the subject matter discussed in the manuscript

Authors and Affiliations

- **Marco Anelli,** MD - Senior Medical Advisor at ProductLife Group, Milano, Italy; MAnelli@productlife-group.com
- **Xavier Fournie*,** MD – Medical Affairs at Icon Commercialisation & Outcomes, Lyon, France; Xavier.Fournie@iconplc.com
- **Zurab Koberidze,** MD, PhD, MPH – Director Pharmacovigilance at FGK, Munich, Germany; zurab.koberidze@fgk-pv.com
- **Flemming Kjaer Jorgensen**; MSc Pharm **–** Pharmacovigilance at Klifo, Glostrup, Denmark; flemming.kjaer.jorgensen@klifo.com
- **Nicolas Tsiakkas, MD –** Scientific Director, Medwork, Athens, Greece; ntsiakkas@medwork.gr
- **Miranda Dollen*,** BSc – Pharmacovigilance at MZD Consulting and Icon Commercialisation & Outcomes, Wimborne, UK**:** Miranda.Dollen@iconplc.com
- **Errietta Economou*,** MSc **–** Head of Pharmacovigilance at Creative, Athens, Greece; eeconomou@creativephs.com
- **Essam Ghanem*,** MD **–** Pharmacovigilance at Vigi-Care BVBA, Gent, Belgium; info@vigi-care.com

All listed authors were members of the **EUCROF PharmacoVigilance Working Group** at the time of initial writing of this paper.

*: Affiliations and contact details may have changed in 2021 for these authors.

The initial version of this paper was endorsed by the EUropean CRO Federation (EUCROF) board in July 2019. The paper was reviewed and updated in September 2021

Correspondence: M.Anelli@productlife-group.com

Table of Contents:

Abstract

There seems to be a lack of clarity about the meaning of terms such as “artificial intelligence (AI)”, “machine learning (ML)”, “data science (DS)” and “office automation (OA)”, which should not be used interchangeably. Pharmacovigilance (PV) represents a very interesting field in this regard, since it poses some unique technical challenges.

For example, the “signal to noise ratio” and the “quality of the information” are very different in reports from clinical trials and in information coming from social networks, smartphone apps or smartwatches, and therefore require different approaches. Case processing, moreover, requires a mix of administrative and repetitive tasks (e.g. data cleaning and form filling) and tasks that require a high level of experience and specialization (e.g. medical reviews and signal detection).

Technology now offers a wide spectrum of solutions, which go from already available “simple” procedure automation allowing significant improvements in efficiency and quality to sophisticated Natural Language Processing (NLP) solutions, which have shown interesting results but are not yet fully operational. A sensible approach, from the perspective of a PV executive, could be to start implementing and reaping the benefits of already available solutions while at the same time keeping the landscape monitored for new developments.

This is the first part of a two-part article. In the second part we will address some of the open issues, of organizational, social and regulatory relevance, concerning the implementation of AI- and ML-based technologies in our domain.

A Widespread Phenomenon

In the field of medicine, the MeSH (controlled medical vocabulary of PubMed) has the term "artificial intelligence" as a descriptor, further subdivided into 13 subheadings. The whole database contains over 130,000 references to Artificial Intelligence (AI).

Already in December 2018, Professor Eric J. Topol of the Scripps Institute hailed AI as "one of the top advances in 2018 that are shaping medicine" (Topol, 2018).

In April 2021, a very important joint workshop was held jointly by EMA and HMA, dedicated to the regulation on AI (EMA, 2021).

Therefore, it is quite safe to assume that artificial intelligence-based applications are going to affect all levels of PV, even though the present situation is still far from being clear.

Both Machine Learning (ML) and Artificial Intelligence (AI) rely heavily on concepts originating from Data Science.

Data science (DS), is an interdisciplinary field that uses scientific methods, processes, algorithms and systems to extract knowledge and insights from data in various forms, both structured and unstructured (Dhar, 2013; Leek, 2013).

DS employs techniques and theories drawn from many fields within the context of mathematics, statistics, information science, and computer science. The term, coined in the 1990s, has recently become more and more important with the advent of "Big Data", another often-overused term. The Wikipedia (Wikipedia Big Data, 2020) definition of Big Data includes the following statement: "Data sets that are too large or complex for traditional data-processing applications and software to adequately deal with" (Snijders et al, 2012).

Big Data, whenever mentioned, is usually associated with its three "key concepts": volume, variety, and velocity (also known as "the 3 Vs") (Laney, 2001). It is not uncommon for a Big Data "spreadsheet" to have up to a few million rows and over a few thousand columns. This may seem a huge amount of data, but such a set can be easily generated by a source such as, for example, a health claims database. Although latest spreadsheet software can handle such volume of data, it would be unthinkable to even start dealing with tables of similar size manually.

The main aim of statistics so far has been to draw meaningful insights and inferences from samples that – ideally – had to be as small as feasible, because of the difficulty and cost of data collection. With big datasets the most complex and challenging task is to consolidate data and reduce the number of variables down to a manageable size, without losing relevant information.

Another possible approach, known as "Data Mining", is to process these vast amounts of data to "mine", i.e. to isolate, only the relevant information.

In conclusion, it can be said that DS is used to get "insights" from datasets, especially large ones.

Machine Learning

Another area where there is sometimes lack of clarity is in the distinction between AI and ML.

Based on the very definition of the two terms, it is safe to assume that the vast majority of PV related applications are not AI-based, but rather ML-based.

A working definition of Machine learning (ML) could be: “The scientific study of algorithms and statistical models that computer systems use to perform a specific task without using explicit instructions, relying on patterns and inference instead” (Wikipedia Machine Learning, 2020).

ML algorithms “build” a mathematical model out of sample data, usually known as "training data", in order to make predictions or decisions without being explicitly programmed to perform the task (Bishop, 2006). An example of ML-based system is a spam filter that, by “observing” which messages you delete and which messages you keep, can progressively improve its accuracy in classifying new messages as either spam or non-spam.

This is also an example of what is called “supervised learning”, in which the machine is confronted with a set of inputs associated with the correct answers (training set). It is then up to the algorithms used by the machine, during the training/learning phase, to try to find what relevant and irrelevant messages have in common and make up the “rules” that will be used to predict whether a mail is or not spam (Guzella and Caminhas, 2009).

Another example of ML is the targeted ads that we receive every time we use search engine or shop online. These systems track our activities and “learn” from them, with the objective to propose goods and services that fit our needs and tastes. The same approach is used to teach computers to recognize faces, patterns, road signs, read x-rays or ECGs, etc.

The training period of an algorithm requires a significant amount of time and resources (both human and machines) and a good set of training data, similarly to what occurs during the training of human beings (Shalev-Shwartz and Ben-David, 2014).

Another possible approach is called “unsupervised learning” and lets the program “discover” analogies and dissimilarities in the input set on its own. The unsupervised approach is inherently and theoretically much more complex. Presently, it can be considered as essentially experimental, except in some data mining applications (Bengio et al, 2012; Längkvist and Loutfi, 2014).

In a very interesting application, an unsupervised system was fed all the PubMed abstracts on vitiligo published from 2012 to 2014 and was able to originate autonomously some relevant “considerations” on the subject, such as the evolution of treatment after the introduction of phototherapy, and the importance of some concomitant conditions (Rani et al, 2015).

Even more experimental is “reinforcement learning”, where the system is fed back the results of its decisions and adjusts its strategies accordingly (Silver et al, 2018). This is the approach used by chess or Go programs like Alphazero, which, like human beings, get better and better the more they play.

ML, in conclusion, is used to make “predictions” based on the available data.

The most frequently used algorithms are based on statistics, while others mimic genetic evolution or the behaviour of neurons (neural networks). “Neural learning” is also known as “Deep Learning” (Sengupta et al, 2019).

The selection and implementation of the most appropriate algorithm, out of the hundreds available, requires in depth domain knowledge and can be a lengthy process.

Another important point to be considered in the selection of an algorithm for operational use is its “documentability”.

Some algorithms, in fact, can be considered as “black boxes”, a technical term meaning that there is no way of telling how the system reached a given conclusion. Neural networks, after a few training sessions, can be very accurate in their predictions, but are essentially black boxes.

Other algorithms, on the other hand, are much more transparent and let the developers know what their “thought processes” were. Decision trees are a good example of transparent algorithms.

This distinction became important when banks started granting or denying credit based using ML algorithms. When some clients, who had been judged at high risk of insolvency sued them, the organizations that were using for example neural network algorithms, were unable to explain in court why the decision was taken and found themselves in a difficult situation (White & Case LLP, 2017; Douglas, 2019).

Artificial Intelligence

Computer science defines AI research as the theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages. AI also involves the study of "intelligent agents", that is, any device that perceives its environment and takes actions that maximize its chance of successfully achieving its goals (Poole et al, 1998; Russel and Norvig, 2003; Nilsson, 1998; Legg and Hutter, 2007).

AI, sometimes referred to as “computational intelligence”, is extensively used in fields such as robotics, game playing, self-driving vehicles, etc.

As previously stated, since AI implies the concepts of decisions and actions, we can assume that most current PV applications belong more to the ML domain.

An example

Doctor David Robinson, Chief Data Scientist at Data Camp, published a very good example that illustrates the differences between DS, ML and AI and how they can be used together (Robinson, 2018).

In detail:

Suppose we were building a self-driving car, and were working on the specific problem of stopping at stop signs. We would need skills drawn from all three of these fields.

- ***Machine Learning**: The car has to recognize a stop sign using its cameras. We construct a dataset of millions of photos of street side objects, and train an algorithm to **predict** which have stop signs in them.*
- ***Artificial Intelligence**: Once our car can recognize stop signs, it needs to **decide** when to take the **action** of applying the brakes. It's dangerous to apply them too early or too late, and we need it to handle varying road conditions (for example, to recognize on a slippery road that it's not slowing down quickly enough), which is a problem of control theory.*
- ***Data Science**: In street tests, which require the analysis of millions or even billions of data points, we find that the car's performance isn't good enough, with some "false negatives" in which it drives right by a stop sign. After analyzing the street test data, we gain the **insight** that the rate of false negatives depends on the time of day: it's more likely to miss a stop sign before sunrise or after sunset. We realize that most of our training data included only objects in full daylight, so we construct, using a data science-based approach, a better dataset including nighttime images and go back to the machine learning step.*

DS, ML and AI therefore should not be used interchangeably.

Automation: From VBA to Programming Environments and Applications

Not all users know that Microsoft Office comes with Visual Basic for Applications (VBA) bundled for free.

VBA is a very complex and versatile programming environment that can access all the functions of Windows, including input and output, and is especially designed to interact with data in Office format (Word, Excel, and Access).

VBA can therefore be used to automate many lengthy and repetitive tasks, which are – by definition – unpleasant to many people and prone to error

The very first example of VBA automation is the recording of an Excel or a Word macro. Without leaving the Office environment, however, macros can be "stitched together Lego-style" in a more complex task-solving structure.

The drug usage data from several thousand patients, for example, could be cross-checked against the health resources usage of the same patients (or of a control population) and automatically consolidated in a series of pivot tables and charts.

Any significant deviations and/or signals could be flagged and – if needed – a warning e-mail could be sent to the appropriate recipients. The information could then be summarized in dashboards, which would allow a single executive to monitor several areas at once, and automatically included in Word/PDF reports or in PowerPoint presentations. An additional advantage of this approach would be that each step of the process can be automatically filed and documented for quality verification purposes, audits, inspections or for future re-analyses.

VBA, however, is not the only player in the field. Sometimes, the situation is made more complicated by the potential issue of the development environment.

Very similar capabilities are offered also by open-source environments, like for example LibreOffice. (Wikipedia LibreOffice, 2020)

Most commercial software companies and AI/ML vendors tend to propose ready-made proprietary solutions. On the other hand, most of the advanced research in academia is done using open-source tools, which are available for free to the scientific community.

This approach offers relevant advantages in terms of budget, but has led, presently, to the lack of a standard procedure. Data scientists, in fact, do not usually employ a single tool, like VBA, from the ground up but prefer to use open-source environments like R or Python, which can interface more easily with data in Office format on one side and with more sophisticated AI- and ML-based tools on the other side.

One example of this two-sided approach applied to PV is represented by literature searching, a very resource intensive task for companies that have a high number of drugs/active principles (APIs).

The query of the appropriate databases (PubMed, EMBASE and/or EVDAS) at given intervals, the collection of the results, the preparation of work-sheets for the medical reviewers, the handling of the retrieval of original papers, the reporting, and all the necessary documentation can all be easily implemented by using automation techniques (Fantini, 2016; Fantini, 2019).

The approach, however, could be taken one step further, by adding an AI/ ML-based system capable of identifying and flagging potentially relevant papers (AI check) to the review process. This is technically called a “classification problem”.

This system could work in parallel with a human medical reviewer or, after adequate training and a validation process, even be allowed perform independently.

Another area that could benefit from automation is the “local literature search” required by regulations from Marketing Authorization Holders (MAHs). For publications in French some publishers already offer dedicated services; for most other languages, however, the problem still remains, exacerbated by the lack of standardization of format, publication schedules, and other issues typical of the smaller local journals. Here again, the approach could be initially automation based, with routines and scripts that periodically check all relevant websites for changes and updates. Whenever an update is found, the relevant content could be downloaded and presented to a reviewer for evaluation and, in parallel, included in a ML based evaluation process.

Focus on Pharmacovigilance

A moving target

Quite a few papers have appeared in recent times about the potential usefulness of AI based approaches applied to PV (Murali, 2019; Hussain, 2021; Letinier, 2021; Sessa, 2021; Hauben, 2021).

The first thing that needs to be taken into account is that the kinds of data that need to be processed and analyzed in PV activities are very diverse and vary greatly as far as their "quantity" and "quality" are concerned.

On one end of the "data spectrum", in fact, we have PV data from clinical trials, where all reports are followed up very closely and where all data is (or should be) checked, monitored and verified.

In line of principle, data from Individual Case Safety Reports (ICSRs) should have the same quality of PV data coming from clinical trials, since the same information is recorded. In reality, – as anybody who has done ICSR follow-up will probably confirm – obtaining the right information from the reporter, be it a patient, a caregiver or a Health Care Professional (HCP), is often a complicated task.

In our opinion, the data originated with the help of the new PV apps, which are becoming more and more important in developing countries, can be considered of similar quality or possibly a little higher. In this area some relevant work was done by the WEB-RADR (Web-Recognising Adverse Drug Reactions) Consortium (WEB-RADR Consortium, 2020).

The initial WEB-RADR pilot app was made available in the UK, Netherlands and Croatia and is now deployed also in Zambia and Burkina Faso.

Peer-reviewed Medical Literature Monitoring (MLM) is another essential source of safety information, and the peer-review system should ensure the reliability of the data. However, the available information is usually limited to what is found in the paper.

Social media platforms are a plentiful source of information; however, it is much more difficult to confirm its reliability.

There is a sort of hierarchy in data from social media, where posts in groups dedicated to Health Care Professionals or to patients and their caregivers are potentially more reliable than data posted on Facebook or Twitter, where the information posted has been proven in several cases to be inaccurate, false or even willfully misleading. This point will be discussed more in depth in the second part of this article.

On the other end of the data spectrum, we have data from "wearable devices" (Fitbits, Garmins, iPhones, etc.) that can originate millions or billions of data points about heart rate, sleep/wake patterns, activity and – in some cases – also oxygen saturation or blood sugar.

A potentially useful example of this approach is the Apple's release of their ECG app for the Apple Watch Series 4 in December 2018. It was the first US Food and Drug Administration (FDA)-cleared consumer deep learning algorithm (a subtype of AI as we have seen) that will potentially affect millions of users. Apple projected it to be purchased by 9 million people by the end of this year.

The user's heart rate is tracked at rest and with physical activity, when there is deviation from the expected pattern for that individual an alert is given to the person to record their single-lead ECG for 30 seconds with an immediate interpretation. Needless to say, this approach has sparked a lot of debate and the writing of several papers on the pros and cons of using such devices (L'Hommedieu et al, 2019; Perez et al, 2019; Whelan et al, 2019).

According to some authors, the results obtained so far are quite promising, even though some hurdles need to be overcome (Koshy et al, 2018; Karmn et al, 2019).

An interesting point, however, is that this approach has already entered "social medicine". Smart watches, in fact, were used to monitor and quantify the operational parameters of a community based cardiac rehabilitation program in the UK (Khushal et al, 2019).

Any future meaningful machine-aided approach to PV, therefore, should be able to exploit – and integrate – different sources and their peculiarities.

Moreover, the widespread use of these new sources will imply a substantial re-thinking and re-shaping of many signal detection and management processes within the PV structure of a MAH.

Case processing

PV case processing is a very structured process, in which multiple professionals are involved and where quality and compliance with timelines are of paramount importance.

For this reason, the past few years have seen a significant number of attempts aimed at automating – partially or entirely – the handling of PV cases.

GSK can be considered as one of the pioneers in the exploration of this domain (see section on social networks), but quite a few other big players have entered the field.

Bayer, for example, has partnered with Genpact to apply AI to case assessment (Sandle, 2018), while Celgene has launched an initiative named Chrysalis, with the objective to: "*exploit artificial intelligence, machine learning and intelligent automation to increase operational efficiency, consistency, quality of data collection, and signal detection*" (Celgene, 2018).

Celgene, who partnered with IBM Watson to perform research activities in this area, has also made some relevant contributions to the research on both the technical and organizational aspects of the issue (Abatemarco et al, 2018; Danysz et al, 2019), described more in detail in the following paragraphs. The results of this approach are promising. A paper published by IBM and Celgene on their joint initiative reported accuracy rates of 98% in some case processing related activities such as WHO-DD coding, or confirmation of the "validity" of a report (Abatemarco et al, 2018).

On the other hand, in other crucial areas, such as adverse event detection from source documents, the accuracy was barely above 75%, a remarkable feat, but still unusable in a real-life situation.

The present trend appears to be toward a probably more attainable goal, such as the "automation" of specific PV related tasks, vs. the creation of a complete end-to-end process.

Recent published examples concern the evaluation of the severity of ADRs (Chauvet, 2020) or of their causality, using innovative approaches such as "knowledge mining" (Bresso, 2021) or "revamping" old approaches such as the tried and true Roussel UCLAF algorithm for drug induced liver injury (Teschke, 2012).

Structured or unstructured?

The main problem that needs to be solved in machine-aided PV case processing, like in many other areas where natural language is involved, is the fact that the data to process can be both structured and unstructured and therefore needs to be processed and analyzed differently.

An example of very structured data is an Access file, where each field has a label and a very specific type of content (characters, dates, numbers, etc.). A CIOMS form can be considered as a partially successful attempt at structuring PV data to facilitate processing.

On the other hand, a narrative is a much unstructured source, where the information is "hidden" in the prose written by a human being and there is no pre-set order or sequence.

The use of predefined formats and/or procedures for narrative preparation is sometime of help, but the very fact that information in natural language is structured in "sentences" is a difficult challenge for all automated approaches.

A Tweet could be considered even more unstructured, with jargon, misspellings, etc. Let's consider, for example a sentence like: *"@drug: OMG, don't take it, my aunt took it last year and died!"*, which would undoubtedly draw the attention of any human reviewer, but would constitute a relevant yet potentially unsolvable problem for any algorithm.

The more unstructured the data, the more "intelligence" is required to process it in a meaningful way.

Natural Language Processing

Natural language processing (NLP) is a subfield of computer science, information engineering and artificial intelligence concerned with the interactions between computers and human (natural) languages, in particular how to program computers to process and analyze large amounts of natural language data.

Sometimes NLP is used in conjunction with "text mining", which applies data mining techniques to unstructured text to give it some usable structure.

According to some Authors : "The main benefit of NLP is in the time savings associated with automation of various medication safety tasks such as the medication reconciliation process facilitated by computers, as well as the potential for near-real-time identification of adverse events for post marketing surveillance such as those posted on social media that would otherwise go unanalyzed" (Wong et al, 2018).

However, a caveat remains: "NLP is limited by a lack of data sharing between health care organizations due to insufficient interoperability capabilities (i.e. communication among the different data repositories and organizations), inhibiting large-scale adverse event monitoring across populations."

From text to MedDRA (and beyond)

A very interesting example of machine-aided case processing has been published by Celgene, which devised a modular "consortium" approach (Abatemarco et al, 2018). Each narrative is submitted to ten separate subsystems, called "cognitive services", whose task is to deal with a specific aspect of the problem, such as the identification of the involved drugs, the evaluation of the seriousness or expectedness, the attribution of MedDRA/WHO-DD codes, etc. The advantage of this approach is that the training, optimization and implementation of each service can proceed in parallel, with an increase in the overall performance of the system every time one service becomes operational.

The effort required by the setup and the training (based on 20,000 cases) has been extensive but – at the moment of the publication of the data – not all "services" had reached a level of accuracy that would allow their use in a production environment.

Follow-up and scheduling

While automated analysis of narratives represents the Holy Grail of machine-aided pharmacovigilance, there are several other resource-consuming and time critical tasks than can be tackled using already available automation-based approaches. By using simple tools like VBA or more advanced development environments, the status of each case can be followed and, whenever necessary, warnings and reminders can be scheduled and issued automatically, including escalations if needed.

As already mentioned, one advantage of this automated approach is that all the steps taken can be automatically documented in detail for quality, review or inspection purposes.

Automated dash-boarding and reporting

Another area that can benefit from already existing and operational technologies is the transformation of raw data into summary tables/charts and the inclusion of these summary representations into reports, dashboards or presentations. The solutions can be implemented using VBA (again), which allows tying together Excel and Word documents or using the tools offered by many commercial vendors.

However, the concomitant use of data mining/ML and AI technologies allows also for the implementation of very sophisticated solutions.

Translations and translation memory

Translations are often a relevant problem for Regulatory Affairs and PV groups: let's consider as an example the handling of narratives or the translation of SmPCs or Package Inserts. Unfortunately, completely automated and unsupervised translation is not available yet. However, some commercially available tools, which employ existing technologies, can help manage both simple and complex translations by first splitting a given text into paragraphs and then proposing an automated translation for each paragraph (Lingohub, 2020).

The status of each paragraph (proposed, checked, finalized, etc.) is tracked by the system and, at the end of the project, the final document is automatically re-assembled, preserving the original format (font, size, tables, etc.). In our opinion, in this situation, the translator becomes essentially an editor/supervisor and the time needed is drastically reduced.

The initial quality of the translation is surprisingly good and the systems, using ML based technologies, can automatically learn how to translate a given expression after a couple of occurrences.

After some "practice", therefore, the system develops a translation memory (TM) database that stores segments, which can be sentences, paragraphs or sentence-like units (headings, titles or elements in a list) that have previously been translated, in order to aid human translators. The TM database can be shared among different translators, assuring a higher level of consistency, which in a domain like PV, can be a relevant bonus.

Other commercial systems offer a "batch" approach, in which files are sent to a cloud-based processing system, which translates them automatically, maintaining the original formatting.

In many cases text mining algorithms (employing variable degrees of "learning" and "intelligence") can also be used to further increase consistency control. A very simple example of this additional automated Quality Control is making sure that – going from one language to another or from one version to another – the units of measure and the figures remain the same. This may not seem a crucial point but, in quite a few cases, an error of this kind in a Package Insert has led to the recall of a batch of product.

Signal detection and management

The EMA website states that "safety signals can be detected from a wide range of sources, such as spontaneous reports, clinical studies and scientific literature. The EudraVigilance database is an important source of information on suspected adverse reactions and signals."

An ideal system should therefore be able to handle presently required sources and outputs but also be able to accommodate potential future developments without disruptions in the work-flow.

Setting up and exploiting the potential of such a system does require the experience and personal skills of PV people. These professional figures are still essential and will not be replaced in the foreseeable future by any form of AI. What machine-aided systems can do is to take care of the most repetitive and time-consuming parts of the work and report/file the results in a way that will let experts concentrate on the important information. Here again, what was described in the "Automated dash-boarding and reporting" section can be applied to the output originated by signal detection and management systems.

In this regard we could quote the words of the Celgene/IBM group who, in the conclusions of their paper, state: "*Through this supported decision-making, pharmacovigilance professionals may have more time to* *apply their knowledge in assessing the case rather than spending it in performing transactional tasks to simply capture the pertinent data within a safety database*" (Abatemarco et al, 2018).

Signal detection and management solutions can be implemented again using proprietary software (such as SAS or SPSS), VBA or more advanced tools.

In this regard, many researchers consider R as a preferential environment for developing PV applications, because of the availability of very advanced statistical and reporting "libraries".

Conclusions

AI and ML based technologies are "here to stay" and will influence several aspects of our daily lives.

The situation is not completely clear at the moment, first of all because the most advanced technologies are still in their infancy and, secondly, because the regulatory framework – which constitutes an essential part of the PV environment – has not been set yet and this does not appear as an easy task.

Several organizations (Agencies, pharma companies, software developers, and academia) have tried to apply Data Science, Machine Learning, Artificial Intelligence or "plain automation" to PV processes, with variable degrees of success.

Presently, an "end-to-end" operationally feasible solution, which "could make PV people obsolete" is not available and is not foreseeable in the near future.

There are, however, some success stories in which sub-processes in the workflow of PV activities have been automated and integrated into the existing flows, with significant increases in efficacy and efficiency.

It can be expected that, with the continuous improvement of technical means and with the "field use" of technologies until now employed only in academia and basic research (ML and AI algorithms are a clear example of this), more and more PV related activities will become "automatable".

A sensible approach, from the perspective of a PV executive, could be to start implementing and reaping the benefits of the already available solutions in his/her organization – and therefore start reaping the benefits – while at the same time keeping the landscape monitored for new developments. In the second part of this article we will examine some of the organizational, social, and regulatory pending issues posed by the implementation of AI and ML based technologies in our domain.

References

Abatemarco D, et al. Training Augmented Intelligent Capabilities for Pharmacovigilance: Applying Deep-learning Approaches to Individual Case Safety Report Processing. *Pharmaceutical Medicine* 2018; **32**(6):391–401.

Abatemarco D., et al. Training Augmented Intelligent Capabilities for Pharmacovigilance: Applying Deep-learning Approaches to Individual Case Safety Report Processing *Pharmaceutical Medicine* 2018:**32**(6):391–401.

Bengio Y, Courville A, Vincent P. Unsupervised Feature Learning and Deep Learning: A Review and New Perspectives. arXiv:1206.5538v1 [cs.LG] 24 Jun 2012

Bishop CM. Pattern Recognition and Machine Learning. New York, Springer; 2006.

Bousquet C, et al. Semantic Categories and Relations for Modelling Adverse Drug Reactions Towards a Categorical Structure for Pharmacovigilance. *AMIA Annual Symposium Proceedings;* 2008;**6**(61):61-5.

Bresso E et al Investigating ADR mechanisms with Explainable AI: a feasibility study with knowledge graph mining BMC Med Inform Decis Mak. 2021 May 26;21(1):171. doi: 10.1186/s12911-021-01518-6.

Celgene promotional material. 2018. [published online]. Available at: https://media.celgene.com/content/uploads/Chrysalis-Celgene-Newsroom-June-2018-21JUN20181.pdf

Chauvet R et al Classification of the Severity of Adverse Drugs ReactionsStud Health Technol Inform. 2020 Jun 16;270:1227-1228. doi: 10.3233/SHTI200375.

Danysz K, et al. Artificial Intelligence and the Future of the Drug Safety Professional. *Drug Safety* 2019;**42**:4:491-7.

Dhar V. Data Science and Prediction *Communications of the ACM* 2013;**56**(12):64-73.

Douglas J. The Apple Card algo issue: What you need to know about A.I. in everyday life. 2019. Available at: https://www.cnbc.com/2019/11/14/apple-card-algo-affair-and-the-future-of-ai-in-your-everyday-life.html

European Medicines Agency Joint HMA/EMA Workshop on Artificial Intelligence in Medicines Regulation 19 & 20 April 2021 Workshop Report. https://www.ema.europa.eu/en/documents/report/report-joint-hma/ema-workshop-artificial-intelligence-medicines-regulation_en.pdf

Fantini D. Querying PubMed via the easyPubMed package in R. 2016. Available at: http://www.biotechworld.it/bioinf/2016/01/05/querying-pubmed-via-the-easypubmed-package-in-r/

Fantini D. Retrieving and Processing PubMed Records using easyPubMed. 2019. Available at: https://www.data-pulse.com/projects/Rlibs/vignettes/easyPubMed_01_getting_started.html

Guzella TS, Caminhas WM. A review of machine learning approaches to Spam filtering. *Expert Systems with Applications* 2009;**36**(7):10206-22.

Hussain R Big data, medicines safety and pharmacovigilance. .J Pharm Policy Pract. 2021 Jun 2;14(1):48. doi: 10.1186/s40545-021-00329-4.

Karmen CL, et al. The Clinical Value of Heart Rate Monitoring Using an Apple Watch. *Cardiology Review* 2019;**27**(2):60-2.

Khushal A, et al. Insufficient Exercise Intensity for Clinical Benefit? Monitoring and Quantification of a Community-based Phase III Cardiac Rehabilitation Programme: A United Kingdom Perspective. *Public Library of Science One* 2019;**14**(6):1-12.

Koshy AN, et al. Smart Watches for Heart Rate Assessment in Atrial Arrhythmias *International Journal of Cardiology* 2018;**1**(266):124-7.

Laney, D. 3D data management: Controlling data volume, velocity and variety. META Group Research 2001 Note. 6 (70).

Längkvist M, Loutfi A. A review of unsupervised feature learning and deep learning for time-series modeling. Pattern Recognition Letters 2014;42:11-24.

Leek J. The Key Word in "Data Science" is Not Data, it is Science *Simply Statistics* 2013; [Blog published online 2013] https://simplystatistics.org/2013/12/12/the-key-word-in-data-science-is-not-data-it-is-science/

Legg S, Hutter M. Universal Intelligence: A Definition of Machine Intelligence *Minds & Machines* 2007;**17**(4):391-444.

Létinier L et al Artificial Intelligence for Unstructured Healthcare Data: Application to Coding of Patient Reporting of Adverse Drug Reactions. Clin Pharmacol Ther. 2021 Apr 18. doi: 10.1002/cpt.2266. Online ahead of print.

L'Hommedieu M, L'Hommedieu J, Begay C, Schenone A, Dimitropoulou L, Margolin G, Falk T, Ferrara E, Lerman K, Narayanan S. Lessons Learned: Recommendations For Implementing a Longitudinal Study Using Wearable and Environmental Sensors in a Health Care Organization. *JMIR Mhealth Uhealth* 2019;*7*(12):e13305. doi: 10.2196/13305.

Lingohub. Available at: https://lingohub.com/. Queried in 2020.

Luo Y, et al. Natural Language Processing for EHR-Based Pharmacovigilance: A Structured Review. *Drug Safety* 2017;**40**(11):1075-89.

Murali K et al. Artificial intelligence in pharmacovigilance: Practical utility Indian J Pharmacol. 2019 Nov-Dec; 51(6): 373–376. Published online 2020 Jan 16. doi: 10.4103/ijp.IJP_814_19

Nilsson N J. Artificial Intelligence: A New Synthesis. Morgan Kaufmann, 1998

Perez MV, Mahaffey KW, Hedlin H, Rumsfeld JS, Garcia A, Ferris T, Balasubramanian V, Russo AM, Rajmane A, Cheung L, Hung G, Lee J, Kowey P, Talati N, Nag D, Gummidipundi SE, Beatty A, Hills MT, Desai S, Granger CB, Desai M, Turakhia MP; Apple Heart Study Investigators. Large-Scale Assessment of a Smartwatch to Identify Atrial Fibrillation. *N Engl J Med* 2019;**381**(20):1909-17. doi: 10.1056/NEJMoa1901183.

Poole D, et al. Computational Intelligence: A Logical Approach. New York, Oxford University Press. 1998, p. 1.

Rani J, et al. Pubmed.mine R: an R Package with Text-Mining Algorithms to Analyse PubMed Abstracts. *Journal of Biosciences* 2015;**40**(4):671-82.

Robinson D. What's the Difference Between Data Science, Machine Learning, and Artificial Intelligence? Variance explained [Blog published online 09 January 2018]. Available at: http://varianceexplained.org/r/ds-ml-ai/

Russell SJ, Norvig P. Artificial Intelligence: A Modern Approach. Upper Saddle River, New Jersey: Prentice Hall/Pearson Education, 2003

Sandle T. Bayer Applies Artificial Intelligence to Medical Cases, *Digital Journal*. [published online 8 November 2018]. Available at: http://www.digitaljournal.com/life/health/bayer-applies-artificial-intelligence-to-medical-cases/article/536194

Sengupta S, Basak S, Saikia P, Paul S, Tsalavoutis V, Atiah F, Ravi V, Peters A. A Review of Deep Learning with Special Emphasis on Architectures, Applications and Recent Trends. arXiv.org > cs > arXiv:1905.13294. last revised 9 Jul 2019.

Sessa M et al. Artificial Intelligence in Pharmacoepidemiology: A Systematic Review. Part 2-Comparison of the Performance of Artificial Intelligence and Traditional Pharmacoepidemiological Techniques. Front Pharmacol. 2021 Jan 14;11:568659. doi: 10.3389/fphar.2020.568659.

Shalev-Shwartz S, Ben-David S. Understanding Machine Learning: From Theory to Algorithms. Cambridge University Press, 2014.

Silver D, et al. A general Reinforcement Learning Algorithm that Masters Chess, Shogi, and Go Through Self-play. *Science* 2018;**362**(6419):1140-4.

Snijders C, Matzat U, Reips UD. (2012). 'Big Data': Big gaps of knowledge in the field of Internet. *International Journal of Internet Science* 2012;**7**:1-5.

Teschke R, Danan G. Idiosyncratic Drug-Induced Liver Injury (DILI) and Herb-Induced Liver Injury (HILI): Diagnostic Algorithm Based on the Quantitative Roussel Uclaf Causality Assessment Method (RUCAM). Diagnostics (Basel). 2021 Mar 6;11(3):458. doi: 10.3390/diagnostics11030458.PMID: 33800917

Topol EJ. 'Eric Topol's Top Advances in 2018 That Are Shaping Medicine', [published online 27 December 2018]. *Medscape*, https://www.medscape.com/viewarticle/906844

WEB-RADR Consortium. [published online]. Available at: https://web-radr.eu/ . Queried in 2020.

Whelan ME, Orme MW, Kingsnorth AP, Sherar LB, Denton FL, Esliger DW. Examining the Use of Glucose and Physical Activity Self-Monitoring Technologies in Individuals at Moderate to High Risk of Developing Type 2 Diabetes: Randomized Trial. *JMIR Mhealth Uhealth* 2019;**7**(10):e14195. doi: 10.2196/14195.

White & Case LLP. Algorithms and bias: What lenders need to know. 2017. Available at: https://www.lexology.com/library/detail.aspx?g=c806d996-45c5-4c87-9d8a-a5cce3f8b5ff

Wikipedia. Big Data. Available at: https://en.wikipedia.org/wiki/Big_data. Queried in 2021.

Wikipedia. LibreOffice. Available at: https://en.wikipedia.org/wiki/LibreOffice. Queried in 2020.

Wikipedia. Machine Learning. Available at: https://en.wikipedia.org/wiki/Machine_learning. Queried in 2021.

Wong A, et al. Natural Language Processing and Its Implications for the Future of Medication Safety: A Narrative Review of Recent Advances and Challenges. *Pharmacotherapy* 2018;**38**(8):822-41.

Acknowledgements

We would like to thank Laura Rivella, Intelligence Officer at ProductLife Group, and Silvia Bondi, Intelligence Specialist at ProductLife Group, for their precious help in the revision of the manuscript and in the validation of sources and references.