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

Part 2: When Machine meets Man

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

Artificial Intelligence (AI) is playing a more and more relevant role in our daily and professional lives.

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” are very different in reports from clinical trials and in information coming from social networks, smartphone apps or smartwatches.

Part 1 of this two-part article gave an overview of the different approaches used to tackle the specific challenges posed by PV, from office automation (OA) to Natural Language Processing (NLP).

Recent developments in the domain of PV have shown that the landscape is evolving quite rapidly.

Social networks, for example, went from “*indispensable source of knowledge*” to “*secondary and optional*” in a matter of months.

In addition, PV has some very specific needs and the usual parameters used to gauge the adequacy of a machine learn (ML) based system – such as precision and accuracy – may not be sufficient.

The human factor should not be overlooked, either.

The implementation of automated systems poses some very relevant ethical, management and legal questions that are just beginning to be addressed by stakeholders and regulators.

Examples of these questions are the final liability in case of problems, the presence of incorrect (or wilfully misleading) information and the impact that these new technologies will have on the workforce.

A relevant number of PV activities can benefit from technologies that are already available and operational, with important increases in efficiency and quality.

The most sensible approach, already investigated or adopted by some companies, calls for the re-analysis of company processes and for the gradual modular automation of the ones that could benefit most, essentially the most time-consuming, repetitive and error prone.

The modules could then be “stitched together” to reach (if and when possible) the ideal objective of a fully integrated and automated PV system, which will let PV specialists focus on the most important parts of their job.

The Scenario

In the first part of this paper ("Artificial Intelligence (AI) in Pharmacovigilance: do we really need it…? - Part 1: the Technicalities") we briefly reviewed the technical aspects of applying automation, Data Science (DS), Machine Learning (ML) and Artificial Intelligence (AI) to the different aspects of pharmacovigilance (PV) activities.

The field is very active and new data, new publications and new events appear almost every day.

It can be safely assumed that AI and ML based technologies are "here to stay" and will influence several aspects of both our working and our daily lives.

The situation is not entirely 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 such as Regulatory Agencies, pharma companies, software developers and academia have tried to apply DS, ML, AI or "plain automation" to PV processes, with variable degrees of success.

Presently, an end-to end operationally feasible solution, which "could make PV specialists 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 existing flows, with significant increases in efficacy and efficiency.

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) we can expect more and more "automatable" 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 this second part we will examine the areas in which change is likely to occur, what its consequences could be, and the areas that could be most problematic during the transition from "man based" to "man/machine based" PV; a transition that any PV organization will eventually have to undergo.

A Controversial Example: Social Media

An example of the complexity and variability of the situation, as far as PV is concerned, is represented by the use of Social Media (Facebook, Twitter, forums, etc.) as a source of data on the benefit/risk ratio of drugs.

The DIA EuroMeeting 2017 held in Glasgow, can be considered a watershed in this regard.

During a dedicated session of this meeting, it was emphasized that in two years there were 22 million posts in Facebook and Twitter discussing potential adverse events for over 1000 drugs, while at the same time, the whole FDA pharmacovigilance database (FAERS) contained only 8.6 million reports from 1968 on.

Case histories of severe ADRs or potentially dangerous product defects discovered by monitoring social media were presented, and a GlaxoSmithKline representative emphasized social media as an important source of information for the potential misuse and abuse of drugs, using bupropion as an example (Anderson et al, 2017).

Doctor Phil Tregunno, from MHRA and Project Lead of WEB-RADR (Web-Recognising Adverse Drug Reactions) group, an international group aimed at improving PV through new technology (WEB-RADR Consortium, 2020), illustrated the point of view of Regulatory Authorities.

The interest raised by these presentations confirmed the relevance of the subject for both Marketing Authorization Holders (MAHs) and vendors, and sparked a debate on how to implement and monitor social network signal strategies which, given the sheer volume of data, could only be based on a ML/AI approach.

The subject of extracting PV data from social networks had been dealt with methodologically in several papers on how to process posts from Twitter and Facebook only, published from 2012 on. However, only recently was its relevance from the point of view of real-life usefulness questioned, first by the WEB-RADR group itself and later by other researchers (Brosch, 2017; Caster et al, 2018; Rees et al, 2018).

WEB-RADR analyzed 38 products across multiple manufacturers and was able to assess a dataset of 40,000 tweets reporting 56,000 product event combinations, which was compared to the 613,000 product event combinations in the WHO Vigibase. No matter the size of the analysed dataset, all groups appeared to come to a fairly consistent conclusion: in general, data mining in social media datasets does not reveal new safety signals or observations.

One possible exception could be “real life” use of a drug, including abuse, misuse or off-label use.

At the French Pharmacovigilance Day, held in November 2018, Doctor Tregunno confirmed that as far as signal detection and validation is concerned – that is the final aim of all PV reporting activities – the present recommendations of the WEB-RADR workgroup are:

- The use of social media data should be optional
- The reporting of suspected adverse reactions in the form of individual cases should not be required
- Follow-up with social media users should not be required
- Insights obtained through social media data which contribute to the safety profile of a medicinal product should be included within relevant regulatory procedures associated with the product, including PSURs and Risk Management/Minimisation Plans

In case of an emerging safety issue, this should be notified in writing to the relevant competent authority(ies) of Member State(s) and to the European Agency

In less than two years, therefore, social networks went from being hailed as an innovative an indispensable source of knowledge on drug safety to a secondary and optional source that “*performs poorly and cannot be recommended at the expense of other pharmacovigilance activities*”.

This position was confirmed again in an official WEB-RADR paper (Brosch et al, 2019), published in March 2019 and at the “International Pharmacovigilance Day”, held in Barcelona on the 12th to 23rd June 2019.

However social media still seem a potentially interesting source of information as witnessed by two recent papers that emphasized the accuracy of "crowdsourced" data (Gartland, 2021) and the fact that by using Tweets it was possible to identify some known ADRs that were not present in the FAERS database (Farooq, 2021).

Issues Still to be Solved

As we already mentioned, the field of machine-aided PV is far from being stable and mature and several issues remain.

This means that decision makers should be more careful in evaluating their options and their planning. At the same time, however, the potential benefits in terms of getting a competitive edge are much higher for the successful innovators.

The Technical Side (a matter of confusion)

As far as technology is concerned, it should be emphasized that automation-based solutions (using VBA, proprietary solutions or more specialized environments such as R or Python) aimed at helping PV specialists perform their task in the most effective and efficient way possible are already mature and ready to be implemented.

As far as AI/ML based technologies are concerned, instead, the situation still needs some (possibly substantial) improvement.

In fact, a 75% accuracy rate can be a very important experimental result, but in reality, a reviewer that misclassifies a potential adverse reaction one time out of four would have an unacceptable performance.

In many real-life PV situations, even a 99% (one error in a hundred) or 99.9% (one in a thousand) accuracy rate would not be acceptable.

For example, if your procedure for screening potential terrorists were: "do nothing" (i.e. don't check passengers for weapons), your classification of non-terrorist would be correct more than 99.9999% of the times (true negatives).

On the other hand, just one terrorist going through security with a weapon (false negative) could have disastrous consequences.

The actual performance of an AI/ML based system, therefore, should be evaluated on a case-by-case basis, using a risk-based approach if needed.

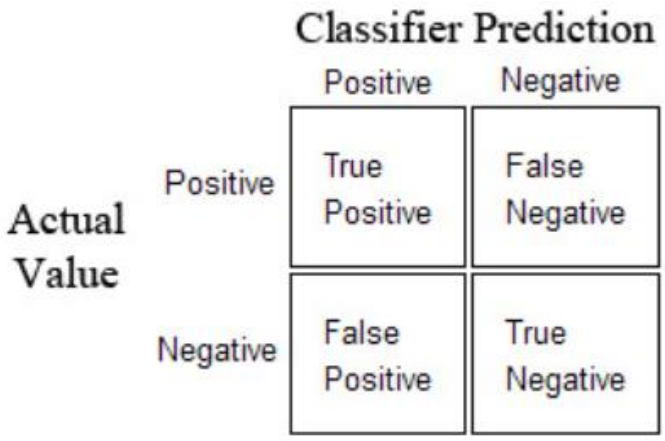

The most widely recognized way to evaluate the performance of a system, and to compare different approaches and solution, is to use a “confusion matrix”.

Similarly to what is done for diagnostic tests, in a confusion matrix the number of correct predictions is compared to the number of errors.

Needless to say, a system that has a somewhat lower accuracy, but tends to err on the conservative side, may be preferable in situations, like PV, where infrequent events should not be missed.

For the moment the best results, from the operational point of view, have been obtained by pooling machine and human resources.

Many systems, in fact, use more than one algorithm on the same piece of information and make a final judgment based on the “decisions” of the single algorithms.

Such a system can also be easily tuned to be more or less strict, according to the number of “votes” required to trigger a positive or negative response. AI/ML systems, as we have seen, require a relevant investment for their set-up and training but, as a rule, their performance tends to improve with “experience”. For these reasons, in some processes, human reviewers and machines are now used in parallel.

This approach has the advantage that, since man and machine use a very different form of “reasoning”, it is more difficult that both will make the same mistake on the same piece of information.

Secondly, when the human and the machine reviewers tend to disagree, the occurrence can be routinely used as a means to improve the performance of the system.

The “Dark Side” (machine against machine)

Already in 2018 a story published by New York Times claims that millions (or tens of millions) of social media accounts are fake, that is, they do not correspond to real people (Confessore et al, 2018).

Many of these accounts are bots, which are computer programs designed to look and behave like a human being.

As we all know, bots have been used in the past to spread fake news about politicians, performers, etc. often successfully.

Therefore, it would not be technically impossible for a company to launch a bot attack against a competitor's drug.

Up to now, we have always considered Health Care Professionals as a trustworthy and reliable source of information.

The fact that now private citizens can also report adverse drug reactions has probably lowered the bar a bit, as anybody who has done PV follow-up work can confirm.

Someday, especially if the situation evolves in the wrong direction, monitoring what is being posted about a drug, or performing identity and reliability checks on reporters, could become a routine part of the follow-up of ADR reports. This is what is routinely done in areas where fake news and smear campaigns are already a (sad) fact of life.

Luckily, (if we want to say so…) "good technology" offers some sort of protection against "bad technology". There are, in fact, AI/ML based text analysis systems that can tell, with a very high degree of accuracy, if a post or a Tweet comes from a bot or from a genuine but disgruntled human being.

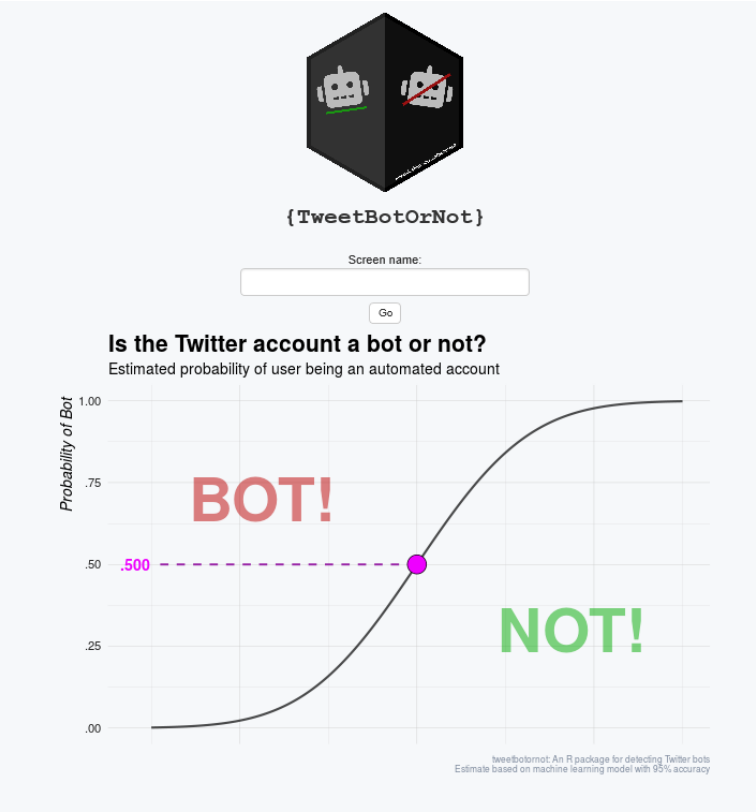

In a series of tests "BotOrNot", now trademarked as "Botometer®" a publicly available ML system showed an accuracy of over 95% in identifying bots vs. real human beings (Sayyadiharikandeh, 2020; Yang, 2020).

Therefore, in a hypothetical PV workflow, we may need to include (and document) a "humanity/reliability check", that is, some sort of step in which, with the help of computer algorithm if needed, the probability of the report coming from a human reliable source is established.

It is safe to assume that the importance of social based reports will remain marginal in the near and middle term future if we sum the difficulty of reliably sifting through millions of posts in non-standard, non-

technical language, the potential for unreliable or misleading information, and the scarce importance placed on social networks as a source of useful and valid PV signals.

The Regulatory Side

As Doctor Tregunno pointed out in his presentations at the French and International PV Days, a more widespread use of machine-aided PV systems (based on simple automation or on ML/AI) will imply a necessary evolution of the existing evaluation, validation and inspection processes.

Since there are no specific regulations pertaining to ML/AI these activities are for the moment regulated by EU 520/2012.

20.6.2012 EN Official Journal of the European Union L 159/5

COMMISSION IMPLEMENTING REGULATION (EU) No 520/2012

of 19 June 2012

on the performance of pharmacovigilance activities provided for in Regulation (EC) No 726/2004 of the European Parliament and of the Council and Directive 2001/83/EC of the European Parliament and of the Council

(Text with EEA relevance)

According to this regulation, ML - and all other processes used in PV - should first meet the requirements to ensure that all quality objectives are achieved in PV activities.

Secondly, since training and qualifications are needed for human resources, it is likely that in the future inspectors will want to know how a machine has been "trained" and how it "learns". AI and ML processes will therefore form part of the QMS (Quality Management System) and all their aspects should be fully documented and available for inspection. This could be a big conceptual issue in systems based on deep/neural learning, which, as we have seen, inherently represent a set of "black boxes".

Companies will also have to ensure that their Quality personnel, including auditors, will have the necessary skills to implement, monitor and audit processes involving AI and ML.

A similar set of skills will have to be provided to inspectors from Regulatory Authorities and, when evaluating results, "acceptable error rates" (random or systematic) will have to be established, tested and agreed upon.

The level of control of the process and the ease (or difficulty) of implementing modifications or corrective and preventive actions, should also be considered.

According to the experts, validation could prove a challenge as well.

The technology, in fact, is still in its infancy and there are no universally accepted methods of evaluation.

ML is by definition based on cycles and repetitions. Therefore, infrequent or "completely new" situations, such as new "problematic drugs", new terminology, new abbreviations, or new events, could theoretically take an unacceptably long time to be detected and acted upon. This is particularly relevant whenever there is no human oversight expected.

There is also a basic difference between how men and machines make mistakes.

Human mistakes tend to be, on average, random and by the law of large numbers; they tend to "cancel" one another.

As Doctor Tregunno, pointed out at the International Pharmacovigilance Day (Barcelona 12-13 June 2019), machine made errors, caused for instance by a faulty algorithm or by wrong parameters, tend to be systematic and to reinforce themselves, therefore skewing the overall performance of the system.

Another potential issue is that companies will probably want to protect their intellectual property and their "trade secrets".

From the point of view of companies, it should be possible to find a common ground, similarly to what happens during "Chemistry Manufacturing and Control" inspections, where inspectors are "exposed" to proprietary and valuable knowledge. Third party vendors and software houses, on the other hand, may be less willing to disclose this kind of information.

The situation clearly calls for a review of the set of rules used and of the way they are managed and applied. Furthermore, an open and transparent collaboration between all involved parties (MAHs, Regulatory Authorities, vendors, etc.) will be needed.

From this point of view, the approach of the MHRA is particularly commendable.

This Agency is in fact already applying internally advanced computer-based technologies (automated follow-up, automated routing, signal detection and management, training) to improve its "incident processing capability" and to be a leading international player in this very important area even after Brexit.

According to some Authors, moreover, while the underlying validation requirements largely remain the same and could draw from existing GMPs, additional activities tailored to intelligent automation are needed to document evidence that the system is fit for purpose.

This is especially important since there are different solutions possible and in use (from rule-based systems to dynamic AI-based systems) and each category needs a unique validation approach (Huysentruyt, 2021).

The need for guidelines and regulation was emphasized also in a wide Italian survey (Stagi, 2020) and, as mentioned in the first part of this paper, in April 2021 a "Joint HMA/EMA Workshop on Artificial Intelligence in Medicines Regulation" took place (EMA, 2021).

The Ethical Side

Elon Musk, the founder of Tesla, considers AI far more dangerous than nuclear weapons. He believes that AI could constitute a threat for the survival of mankind and that stricter government regulations on the subject are necessary (Elon Musk, 2018).

In this regard, he quoted the fact that AlphaGo, an AI based program, taught itself how to play Go without the need of any human supervision and, in a few days, went far beyond the capabilities of any human Go grandmaster (Vincent, 2017).

Musk's concern is that an unregulated AI may quickly develop beyond human control and become unpredictable.

This debate involves much wider areas than PV, like for example automated driving (and the related legal liabilities) or AI based weaponry, and is therefore beyond the scope of this article.

The issue is nevertheless being investigated; in April 2019 the EU published the Ethics Guidelines for trustworthy AI (EMA, 2019).

According to the guidelines, trustworthy AIs should be:

- lawful - respecting all applicable laws and regulations
- ethical - respecting ethical principles and values
- robust - both from a technical perspective while taking into account its social environment

This may sound a bit obvious, however the guidelines put forward a set of 7 key requirements that AI systems should meet in order to be deemed "trustworthy".

A specific assessment list also aims to help verify the application of each of the key requirements.

In addition, a piloting process will be set up to gather practical feedback on how the assessment list can be improved.

All interested stakeholders can already register their interest to participate in the piloting process that has been kicked-off in summer 2019.

A forum discussion was also set up to foster the exchange of best practices on the implementation of Trustworthy AI.

The ethical aspects of PV-related AI can be considered as subsets of the ethical aspects of AI in Medicine.

The debate is strong in this area as well, with papers on the subject published throughout the world, from the JAMA to Chinese journals (Emanuel et al, 2019; Zhang, 2019).

Given Hippocrates's rule "*Primum: non nocere*", everybody seems in agreement on the fact that patient safety should come first, especially when children and disadvantaged categories are involved.

A risk-based approach is also deemed necessary, with case by case "contextualization" of what theoretical parameters such as precision and accuracy may mean when applied to the situation of each patient.

In many situations in fact, such as for example the diagnosis of a treatable serious condition, false negatives are unacceptable, while in other cases it could be the other way around.

In this regard, at a conference held in Boston on "Precision Medicine", Jonathan Zittrian, a professor at Harvard Law School, provocatively stated: "*I think of machine learning kind of as asbestos. It turns out that it's all over the place, even though at no point did you explicitly install it, and it has possibly some latent bad effects that you might regret later, after it's already too hard to get it all out*" (Ross, 2019).

At the same conference another important point was discussed: systematic errors and bias. This point was very similar to the one raised by Doctor Tregunno in Barcelona.

There is some published evidence suggesting that algorithms used to identify skin cancer may be less effective in people with darker skin, thus leading to potentially significant "race based" disparities in diagnosis and treatment (Adamson et al, 2018).

An even greater risk is posed by the fact that the data sets used to train the algorithms may themselves be biased, leading to a possible skewing of responses.

This possibility was pointed out as early as 2008, when a team of Israeli researchers reported that: "*a model trained explicitly to capture a specific operational performance criteria such as attractiveness rating, implicitly and concomitantly captured basic human psychophysical biases and demonstrated a wide range of human-level characteristics of facial attractiveness judgment*" (Kagian et al, 2008).

If the principles are clear, it is probably still too early to foresee if and how they will be implemented in clinical practice, PV practice and/or in regulations.

The Human Side

The question many PV people are asking is: "Will AI/ML become the next *off-shoring*?"

Many EU and American PV professionals, especially at the lower expertise levels, have already lost their job because of the competition coming from countries where labour costs are much lower.

During the Industrial Revolution, machine substituted men (and animals) in those processes where strength and speed were needed, while a similar phenomenon took place more recently when many manufacturing activities were taken over by robots.

Some sort of employee uneasiness in this regard is therefore quite understandable.

Authoritative sources such as Forbes state that AI doesn't eliminate jobs, it creates them. Other sources state instead that the development of AI will make customer service representatives, drivers, factory workers, administrative workers and even computer programmers obsolete.

Finally, there are publications such as The Guardian and the Smithsonian that recall the Luddites of the 19th century and state that the "rebellion" of the Unabomber (an *ante litteram* American terrorist who from 1978 to 1995 kept mailed bombs to academics and technology experts) was originated by his reaction to a "technological change that was destroying human civilization, ushering in a period of dehumanized tyranny and control".

Through the teachings of economy and history we can infer that, if widespread use of AI/ML causes an increase in efficiency and cost/effectiveness in any productive process, chances are that – eventually – such widespread use will take place.

This phenomenon will not be limited to PV and it is difficult to predict whether it will cause major social changes.

However, it can be assumed that PV workers, like many other professionals, will need additional sets of skills.

In addition to a scientific background, AI and ML literacy will probably be needed, more or less in the same way that today basic computer literacy is a must and advanced skills represent a big advantage in practically any field.

The “career path” in PV will probably also have to be rethought.

Nowadays, it is quite common to start “at the bottom of the ladder” doing data entry and data cleaning.

This is usually followed by follow-up and data analysis tasks, report preparation and, for the most skilled ones, by coordination and management jobs.

In a situation where a large part of data entry, case processing and reporting are automated tasks, it will be necessary to find other ways to let junior people gather a sufficient amount of field experience.

Finally, in 2012, the Harvard Business Review wrote that Data Scientist was going to be “The Sexiest Job of the 21st Century” (Davenport et al, 2012).

This prediction has been largely confirmed in the following years and there is a good chance that the “sexiest jobs” of the next decade or so will be “interface jobs”, that is jobs where Data Science, ML and AI interface with science, medicine and PV and where specialists who can “speak both languages” will be in great demand.

The Reward: From Data to Intelligence

In the past few years, the trend has been going towards a deeper involvement of PV in many strategic company activities and it is unlikely that this will change.

It has been discovered that PV is no longer an “unproductive, necessary evil”, but that it can have a significant impact on the “bottom line” of most Pharma companies.

This has prompted some experts to state, perhaps a bit provocatively, that many organizations could benefit from a “Chief PV Officer”.

For example, it is quite common, nowadays, to find a representative from PV in due diligence teams.

In fact, especially in a regulatory situation where referrals and re-evaluations of the benefit/risk ratio of a drug are becoming more common, relevant gaps in the PV documentation could negatively affect the market value of a drug (or a company) considered for acquisition.

The benefit/risk ratio profile of a drug, especially if compared to its competitors, is also very important at the time of the filing of a dossier for pricing and reimbursement.

This has shifted the role of PV from a “reactive” to “proactive” one, that is, from simple “event counting” to the collection, analysis and evaluation of “PV intelligence” in whole therapeutic areas. This implies, within the scope of any PV group, the need to set up and implement new processes and to integrate them with the existing activities.

In fact, before deciding on an investment, many companies now routinely “profile” their new or old drugs and their competitors in terms of efficacy, safety and real-life use, and use this knowledge to support their strategies.

In this profiling activity, sources such as social networks – even if not presently considered relevant by PV regulations – represent an extremely useful source of information and knowledge.

A widespread off-label use could be both a problem and an opportunity, while the potential for abuse or misuse is definitely a risk that needs to be recognized, evaluated and accounted for.

With the use of the appropriate automation technologies, the preparation of a "smart profile integrating different sources such as social networks, scientific literature, lay-press, institutional websites, etc. is relatively easy.

These profiles could be prepared just once, in order to respond to a one-time strategic question or, with very little additional effort, represent a way to continuously monitor specific areas of interest, in conjunction with other company functions like Medical, Marketing or Sales.

In fact, what does not technically constitute a "signal" for PV purposes, may still be extremely useful – if properly exploited – to prevent uninformed decisions or to gain a competitive edge.

Modularity: the Best Approach?

As we have seen, the technology, especially as far as ML and AI are concerned, is not mature yet and a rapid rate of change can be expected for at least a few years.

A single "off the shelf" package that will solve all PV problems is still quite far in the future, but the available technologies, in their present state, can already offer significant improvement in terms of resources and quality.

The "Celgene" approach probably represents the most sensible approach at the moment.

Instead of looking for a single "one shot" solution, it would be better for a company to re-analyse the internal processes and work-flows to see which parts could be successfully automated right now, and which parts could be considered for the near future and which should be left to human beings.

The development and the operational implementation of each "module" could proceed in parallel, compatibly with the allocated resources, and would gradually increase the efficiency and productivity of the system with the minimum possible disruption of the existing flow.

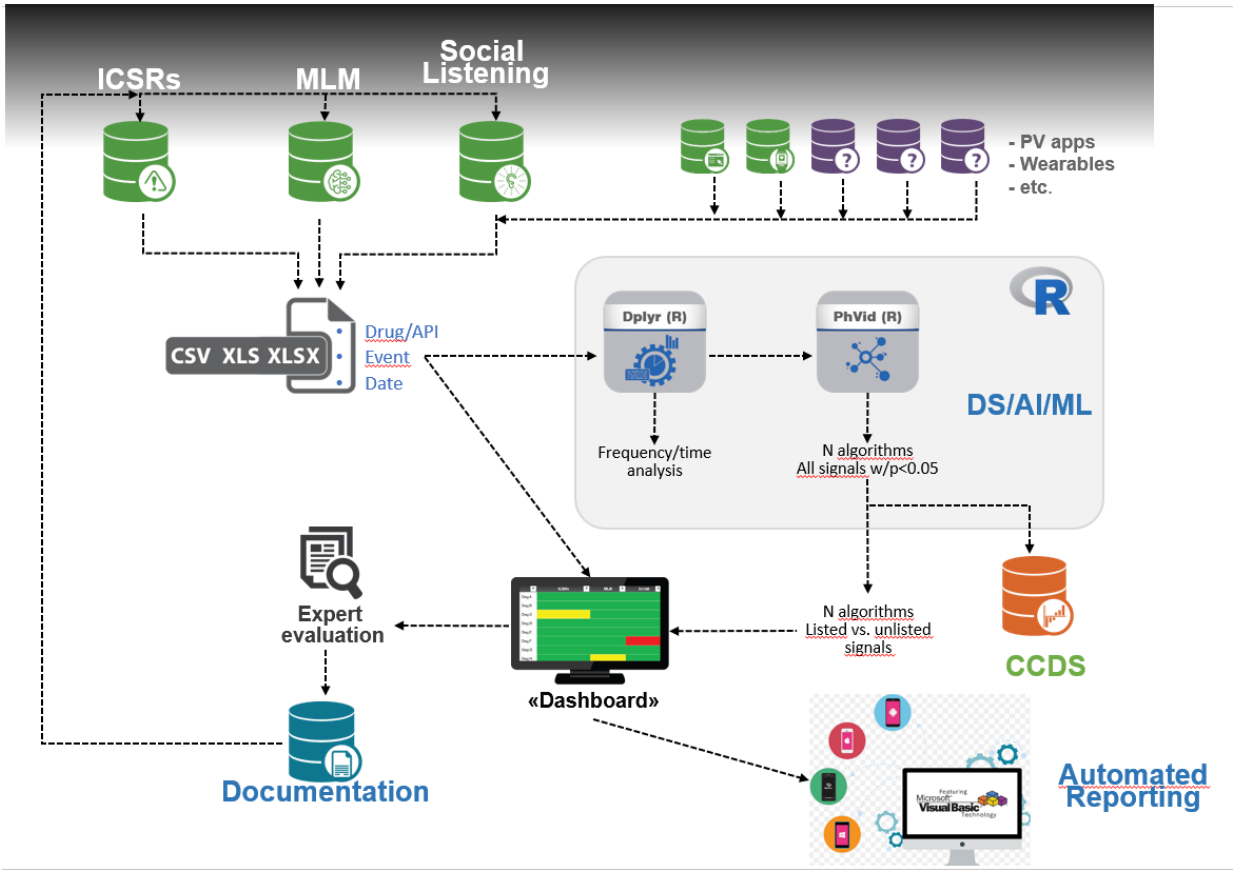

A typical machine-aided PV modular system allows the analysis and integration of all existing structured or unstructured sources, considered together or separately. This is important since some information sources (e.g. social networks) may be relevant for an internal report, but not for an official document like a PSUR.

The modular structure also allows the integration – when needed – of other sources that may become useful or mandatory in the future.

The common format to exchange data should be an industry universal standard, such as Microsoft Office or CSV (Comma Separated Values) for very large datasets.

The integrated information can be used directly to pilot a dashboard or to generate different forms of reporting (in a standard exchangeable format). It can also be fed into more advanced ML/AI based modules that will provide additional analyses, exploiting the existing knowledge base of the company (for example the Core Company Data Sheets or the results of past literature searches).

Eventually, the produced content in its entirety will have to be supervised by experts and each step will be automatically documented for quality, re-analysis or inspection purposes.

How About the Future?

It is difficult to foresee possible future developments in a domain that is changing very rapidly.

As far as PV is concerned there will probably be a trend moving towards a higher integration of the different sources of information, along the lines of what the FDA is presently doing with its "Sentinel" initiative (FDA, 2019)

Sentinel's objective is to securely access information from large amounts of electronic healthcare data, such as electronic health records (EHR), insurance claims data and registries, from a diverse group of data partners.

EMA has been working along the same lines with its "HMA-EMA Joint Big Data Task Force" (HMA-EMA, 2019).

Also other authors agree on the fact that: "*future work in this area (medication safety) will focus on the integration of data sources from different domains to improve the ability to identify potential adverse events more quickly and to improve clinical decision support with regard to a patient's estimated risk for specific adverse events at the time of medication prescription or review*" (Wong et al, 2018).

It is also safe to assume that technologies that are already mature or near maturity will be quickly adopted with the aim of improving efficiency and quality.

The widespread adoption of more "experimental" technologies, based on ML and AI, on the other hand, will probably require their "validation", both from the operational – that is, their actual performance – and the regulatory point of view.

Since the regulatory framework has not been fully established yet, the first step will have to be a collaboration effort among the involved parties.

Like the introduction of machines in factories or of PCs in offices, these new technologies will probably change both the way we work and the set of skill needed to operate successfully – at all levels – in the field of PV.

Conclusions

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

It is therefore unrealistic to expect that PV will not be affected by these changes.

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. Some very important ethical and legal issues have to be confronted as well.

It is unclear to what extent the more widespread use of AI and ML will "socially" affect the workforce in PV, especially workers with lower experience levels, who will probably have to become "AI/ML literate" at some point in their career path.

The potential benefits and the competitive edge that can be presently gained by successful innovators in this field, however, are proportional to the present uncertainty and this is the reason why many big players have already entered the field with significant investments.

It is still quite possible, though, that these investments may prove not rewarding, as the case of social networks in PV has shown.

It should be considered that employing advanced digital technologies in a company's PV processes is not an "all or nothing" decision.

A relevant number of PV activities can benefit from technologies that are already available and operational, with important increases in efficiency and quality, others will become operational shortly.

One example of such areas is the mapping of terms to and from the different medical dictionaries (MedDRA, PRO-CTCAE, Sno-Med, etc.), a task the automation of which has already given some very good results in different settings, such as the Celgene/IBM initiative, the WEB-RADR app and the clinical research (Chung et al, 2019).

Some of these technologies can also be used in "strategic PV activities" such as due diligences, determination of the market value of a drug or company, or market access.

The most sensible approach, already investigated or adopted by some companies, calls for the re-analysis of a company processes and for the gradual and modular automation of those processes that could benefit most (e.g. the most time consuming, repetitive and error prone).

The "automated modules", starting from the ones based on already existing and reliable solutions, could be deployed operationally according to the company plans and budget and each successful step would increase both efficiency and quality.

Gradually, the modules could be "stitched together" to reach (if and when possible) the ideal objective of a fully integrated and automated PV system, which will let PV specialists focus on the most important parts of their job.

References

Adamson A. S., et al. 'Machine Learning and Health Care Disparities in Dermatology', JAMA Dermatology 154:11, (November 2018), pp. 1247-1248.

Anderson L. S. et al., 'Using Social Listening Data to Monitor Misuse and Nonmedical Use of Bupropion: A Content Analysis' *Journal of Medical Internet Research Public Health Surveillance* 3:1 (Jan-Mar 2017), pp. 2-41.

Brosch S et al., 'Establishing a Framework for the Use of Social Media in Pharmacovigilance in Europe' *Drug Safety* (2019), pp. 1-10.

Brosch S. 'Implications of WEB-RADR for Pharmacovigilance' International Society of Pharmacovigilance (ISoP) 17th annual meeting, Liverpool, 2017.

Caster O. et al., 'Assessment of the Utility of Social Media for Broad-Ranging Statistical Signal Detection in Pharmacovigilance: Results from the WEB-RADR Project' *Drug Safety*, 41:12 (December 2018), pp. 1355-1369.

Chung E, et al. Patient free text reporting of symptomatic adverse events in cancer clinical research using the National Cancer Institute's Patient-Reported Outcomes version of the Common Terminology Criteria for Adverse Events (PRO-CTCAE)' *Journal of the American Medical Informatics Association* 2019;**26**(4):276-85.

Confessore N. et al. 'The Follower Factory' The New York Times, 27 January 2018. Available at: https://www.nytimes.com/interactive/2018/01/27/technology/social-media-bots.html

Davenport, T. H. et al., 'Data Scientist: The Sexiest Job of the 21st Century.' *Harvard Business Review* 90: 10 (October 2012), pp. 70–76. Available at: https://www.hbs.edu/faculty/Pages/item.aspx?num=43110

Elon Musk's Last Warning About AI, Reddit Storytime Channel, published 2 October 2018. Available at: https://www.youtube.com/watch?v=B-Osn1gMNtw

FDA (US Food and Drugs Administration. 'FDA's Sentinel Initiative'. Available at : https://www.fda.gov/safety/fdas-sentinel-initiative

EMA (European Medicines Agency) 'Ethics Guidelines for Trustworthy AI', 8 April 2019. Available at: https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai

Emanuel E. J. et al. 'Artificial Intelligence in Health CareWill the Value Match the Hype?', *Journal of the American Medical Association*, 321:23 (May 2019), pp. 2281-2282.

Farooq H et al. Leveraging digital media data for pharmacovigilance. AMIA Annu Symp Proc. 2020; 2020: 442–451.

Gartland A et al. Developing Crowdsourced Training Data Sets for Pharmacovigilance Intelligent Automation. Drug Saf. 2021 Mar;44(3):373-382. doi: 10.1007/s40264-020-01028-w. Epub 2020 Dec 22.

HMA-EMA (Heads of Medicines Agency and European Medicines Agency) Joint Big Data Taskforce Summary Report' EMA/105321/2019, (13 February 2019). Available at: https://www.ema.europa.eu/en/documents/minutes/hma/ema-joint-task-force-big-data-summary-report_en.pdf

Huysentruyt K et al Validating Intelligent Automation Systems in Pharmacovigilance: Insights from Good Manufacturing Practices. Drug Saf. 2021 Mar;44(3):261-272. doi: 10.1007/s40264-020-01030-2. Epub 2021 Feb 1.

Kagian A. et al., 'Machine Learning Predictor of Facial Attractiveness Revealing Human-like Psychophysical Biases' *Vision Research*, 48:2 (January 2008), pp. 235-243.

Rees S. et al., 'Using Social Media in Safety Signal Management: is it Reliable?' *Therapeutic Advances in Drug Safety* 9-9(10), (August 2018), pp. 591-599.

Ross C. 'What if AI in health care is the next asbestos?' *STAT,* 19 June 2019. Available at: https://www.statnews.com/2019/06/19/what-if-ai-in-health-care-is-next-asbestos/

Sayyadiharikandeh et al. Detection of Novel Social Bots by Ensembles of Specialized Classifiers. arXiv:2006.06867

Stagi L et al.Pharmacovigilance and the digital world in Italy: presentation of the results of a national survey. Ther Adv Drug Saf. 2021 Feb 10;12:2042098620985991. doi: 10.1177/2042098620985991. eCollection 2021.

Vincent J. 'DeepMind's Go-playing AI Doesn't Need Human Help to Beat us Anymore' *The Verge* 18 October 2017. Available at: https://www.theverge.com/2017/10/18/16495548/deepmind-ai-go-alphago-zero-self-taught

WEB-RADR Consortium. Https://web-radr.eu/

Yang K. Scalable and Generalizable Social Bot Detection through Data Selection Vol. 34 No. 01: AAAI-20 Technical Tracks 1

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* 38:8 (August 2018), pp. 822-841.

Zhang W. G. 'Medical ethic consideration of artificial intelligence medicine' *Zhonghua Nei Ke Za Zhi*, 1:58(6) (June 2019), pp. 412-414.

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.