EVERSANA™

AI AND THE FUTURE OF MIS
ARTIFICIAL INTELLIGENCE AND THE FUTURE OF MEDICAL INFORMATION SERVICES

The Rise and Fall and Rise of AI

Thomas Barton,
Senior Director, Business Process,
Medical Information and Pharmacovigilance, EVERSANA

A few years ago, we witnessed the heralding of the imminent rise of artificial intelligence (AI) to support medical information services provided through industry-based contact centers. And, in some cases, the message Lonnie Corant Jaman Shuka Rashid Lynn (better known as Common) shares in recent Microsoft commercials is our reality: we do have more power at our fingertips than generations before us. But witnessing the rise of AI and actually experiencing it are two different realities. In retrospect, for medical information services these emerging technologies were overhyped when first introduced to the market in terms of how feasibly they could be applied in practical applications capable of creating a positive return on investment. Many larger companies invested in machine learning 'pilot' programs that appeared to produce less than positive outcomes, and the term "over innovating" entered the corporate lexicon.

The gold rush towards AI has been fueled by exaggerating its capabilities and overgeneralizing the ways the technology could be applied in real world settings. Well-meaning experts talked about the amazing potential of AI – from improving patient outcomes, providing better engagement between healthcare providers and patients, and increasing compliance and efficiency – without clarifying the current limitations. Simply, expectations were created that could not possibly be met by the nascent technologies available at the time, and collectively as an industry we got caught up in this frenzied environment – missing the real opportunities inherent in what the technology could actually accomplish.

And yet, regardless of these setbacks, steady progress in the field of AI has been made over the past several years. The technology has matured, slowly but surely, and through incremental improvements has emerged from the ashes of numerous failed initiatives with impressive gains over its overhyped predecessors. There are numerous successes in sectors within the healthcare industry. For example, AI has taken wearables to the next level. You don't need to look any further than your wrist: smart watches and fitness trackers. A recent report by Global Market Insights, covered this past July by WDTNN, shares that the AI empowered wearable market is poised to hit $180B by 2025. Viz.ai is using AI to detect stroke. They are using AI to synchronize stroke care and envision using AI to improve access to life-saving therapies.

At the same time, many manufacturers are cautiously returning to the table to sponsor new projects based on AI. The FDA is considering a total product lifecycle-based regulatory framework for AI and

> "It will be important that any unsupervised AI system that interacts directly with healthcare professionals or consumers has the ability to detect complaints and either process the complaint appropriately or escalate to human agents. Failure in this area could have serious patient safety and regulatory repercussions."

machine learning technologies. To review their proposed framework, download the report.

CHALLENGES TO AI IN PROVISION OF MEDICAL INFORMATION SERVICES

Regulated companies, who tend to be slow adopters to begin with, waited and watched as other industries gambled with AI investments. In truth, even if these emerging technologies had proved to be extremely successful, it is likely that regulated industries would still have hesitated to adopt them simply due to the regulatory and perceptual hurdles involved in implementing new technologies in this space. One of the main challenges with AI powered technologies for FDA regulated companies is a process known as computer system validation (CSV), a combination of risk assessments and software testing documents designed to prove that the software poses no risk to patient safety or quality of care, is fit for use in a regulated setting, and produces information or data that meet a set of predefined requirements.

Computer system validation is centered on the concept of predictability in the behavior of the system being validated. When creating validation test scripts (which are based on the authorized system requirements) the two key columns are the expected results and the actual results. The expected result column details how the system is expected to behave when the tester performs a series of prescribed actions within the system. The actual results describe how the system actually behaved when the prescribed actions were executed. A test step will fail if the expected and actual results are not the same.

For regulated systems, any change to the underlying code or configuration may require full or partial revalidation of the system to ensure the system is still functioning as expected per the authorized system requirements. This traditional method of verifying system behaviors poses a particular challenge for machine learning-based software due to the potentially dynamic nature of the underlying logic. Many AI-based systems are designed to continuously learn as new information is fed into the algorithm, which arguably constitutes a change in the logic that determines the system output. In this respect, a machine learning-based system behaves much like a human agent, and its functionality is no longer hard-coded and entirely predictable. This situation is further exacerbated by the inscrutability of the machine learning algorithms. Reportedly, even experts in AI models do not always fully understand why the machine learning algorithms they create produce certain results.

In addition to the challenges posed by computer system validation requirements, there are several other aspects of operating within a regulated setting that present significant hurdles to the medical information industry in implementing AI in this context:

- **Adverse event and product complaint recognition.** It is common for human agents working within industry-based medical information centers to detect an adverse event or product quality issue during the course of a phone call where the caller did not specifically call to report a complaint. There is, of course, regulatory requirement for the industry to report all identified adverse events or product quality complaints by maintaining a system for the systematic detection, collection, assessment, and monitoring of these reports. It will be important that any unsupervised AI system that interacts directly with healthcare professionals or consumers has the ability to detect complaints and either process the complaint appropriately or escalate to human agents. Failure in this area could have serious patient safety and regulatory repercussions.
- **Unsolicited requests for off-label use information.** The FDA prohibits the promotion of off-label uses for pharmaceutical products. On a practical level, this means that no person (or system) should provide off-label information about a pharmaceutical product that was not explicitly requested in an unsolicited manner. This has sometimes been referred to as using the analogy of walking into a restaurant and ordering a meal without a waitress or a menu. Further,

responses should be handled in keeping with the other principles spelled out in FDA Guidance including one-on-one communication, answering only the specific question asked and responding with information that is truthful, non-misleading, accurate and fairly balanced. The risk with AI-powered systems is that they could misinterpret the question and return off-label information to a requester that was not related to the actual request, which could be construed as promotional activity.

DO THE PUZZLE PIECES FIT?

Despite these formidable challenges, the provision of industry-based medical information services otherwise presents an almost ideal use case for the application of AI. The reason for this is straightforward: the responses to a high percentage of most known or expected inquiries about any given pharmaceutical product are carefully scripted by the Medical Affairs team and then provided to the medical information contact center team for dissemination. This fact greatly reduces the burden on AI since it will not be necessary to rely solely on machine learning to generate the response; rather, the technology can be almost entirely focused on recognizing the question that was asked. Once the question is accurately identified, a standard response can be identified, tailored if needed and triggered.

This question/response pairing approach can be used to mitigate the risks associated with more complex implementations of AI. In many applications of machine learning, huge data sets are needed to 'train' the AI how to respond. This training could involve providing the machine learning algorithm with tens of thousands of medical inquiry response records, or a careful, laborious, and continuous process of training the AI on the product documentation. These more complex approaches to curating AI knowledge place a huge burden on the technology to accurately understand the content that is being consumed and then to generate an acceptable and compliant response in human readable form. In its current form, AI clearly has not mastered the skills necessary to perform these functions in a regulated setting (though certainly that day will come).

And, medical information teams need to keep these questions in mind when weighing the balance:

- What capabilities would AI need to possess before we could trust it to handle adverse events and product complaint intake?
- What are the risks associated with allowing AI to provide off-label information to healthcare providers?
- How do you validate AI?
- What methods can you use to measure the efficacy of automation technology in terms of customer satisfaction and efficiency?

For the time being, a measured, risk-based approach that takes into account the current limitations of AI is what's needed in order to move forward in this space. There are many use cases based solely on question/response pairing that would greatly benefit the industry without introducing undue risks. The day will come when chatbots can be trusted in a regulatory context - companies such as Clinc.com are pushing boundaries in this field. Think of "Alexa" – the technology behind this could be a game changer for healthcare providers and patients. Yet, most successful applications of AI today, in any industry, are centered on back-end processing tasks and augmentation of human capabilities – or so-called "Intelligence augmentation" (IA).

CREATING A TECHNOLOGY ENABLED WORKFORCE

Rather than waiting for the technology to catch up with our dreams of applying emerging technologies, an analysis should be done to ascertain what the technology can accomplish today, in its current form. There is a common saying about AI in the technology forums: "Machine learning is like having a million low-skilled interns and not a single Einstein." This is how we should be thinking about, and applying, AI today. What would you do with a million interns?

This simple progression of AI, from a back-end processing tool to a front-end optimizer, can greatly reduce the repetitive and menial tasks that would otherwise need to be performed by a human agent. Many compliance risks would be reduced since the technology would presumably be more consistent in its output and each response would have a final review and approval step by a human agent prior to release of the information. At the same time, this type of back-end application of AI provides companies with invaluable experience working with the technology, and abundant data to justify the use of AI in increasingly more complex and autonomous ways.

Will AI replace human beings in the workforce shift? No, humans will always be needed, but incrementally, as the technology improves and the industry adapts to its use, there would be a slow but steady transition towards unsupervised AI and fully automated solutions such as chatbots, intelligent voice agents or smart speakers. This transition will allow the highly trained human agents to focus on the more challenging inquiries that AI won't be able to handle in the near term: training and maintaining the automated systems, "rescuing" the AI when needed, providing quality control and oversight, and addressing complex clinical questions. In the long term, the human skills will shift away from menial tasks towards management and content curation of the AI systems.

The future of medical information services definitely includes automation; in fact, AI has the potential to revolutionize the life sciences industries in general.

Although the challenges outlined in this article are not trivial, few professionals in the industry would doubt that within a decade or so most medical information inquiries will be handled by AI. It is not a question of if, but when, and how we will reach that goal.

HOW DO WE MOVE CLOSER TO THIS GOAL WHERE THE PAYOFFS WILL BE HUGE, FOR PHARMA AND PATIENTS ALIKE?

- Build/obtain technology platform that can support innovation
- Find the right technology partners; pursue a "hybrid" solution
- Implement live chat (AI "rescue")
- Initially focus on lower risk applications of AI that involve humans, such as case QC, recommender system, and case templating
- Analyze your historical data. What percentage of inquiries are simple enough for current AI technology to handle?
- Run a pilot – get your hands dirty. Avoid "pilotitus" (use SMART goals, have clear exit points)

We can thank Dennis Gabor, the physicist who won the 1971 Nobel Prize in Physics for developing the theory of holography, for this quote: "The Future cannot be predicted, but futures can be invented." AI has tremendous potential to improve efficiency and regulatory compliance, not just in the provision of industry-based medical information services, but across the healthcare and life sciences industries. It may be 'just what the doctor ordered.'

This article was originally published in the August 2019 issue of *Med Ad News*.