



Image credit: [Copyright: ivengo](#)



Amit Arora
Jul 12, 2021

AI-Driven Drug Safety

How AI Automates Drug Safety Reporting to Government Health Agencies: A Product Manager's Approach to Innovation

Pharmaceutical companies need to keep a close eye on the adverse reactions their drugs can produce in the population. And not just in one country, but everywhere people use its drugs. However, getting a handle on these reactions can be time-consuming, with thousands, if not millions, of data points requiring management. The challenge of picking out the relevant data to drive insight can be monumental.

This article provides experience-based learnings on how to tackle the challenge of converting a vast, continuous, multi-source stream of incoming data on drug reactions into two categories—serious cases and non-serious cases. It then explores how to automate this classification into a periodic report for review by the FDA and regulatory bodies.

The solution design incorporates the application of artificial intelligence (AI) and automation to extract, re-format, classify, and report adverse events (AE) from multiple sources within FDA-stipulated reporting timelines. The solution transformed the pharmacovigilance operations of multiple global pharmaceutical companies.

The payoff for a typical pharmaceutical firm is big, including:

- Reduction of FDA and regulatory penalties in the range of \$8 million to \$12 million per year.
- Savings of about \$5 million per year from efficient, automated detection and case processing of adverse events.
- Additional IT cost savings of about \$1 million per year from reduced spending on cloud, infrastructure, and licensing.

Key Takeaways

- Learn about the process of building an AI-based platform that processes cases by reading text, understanding human speech, extracting and codifying AE data, and continuously learning.
- The platform uses TensorFlow and home-built algorithms to create a contextual map of keywords that better predicts and prevents adverse effects caused by medicines.
- The solution can streamline signal detection and evaluation, enabling the company to be more proactive and responsive.
- The system continually learns from the wider healthcare delivery ecosystem, accelerates and enhances signal detection and evaluation, and better protects patient safety.

Challenge

Pharmaceutical companies constantly collect, analyze, and report suspected adverse events from a

wide range of sources; including observational and clinical studies, consumers, and healthcare professionals. Moreover, regulations vary from country to country, as do reporting requirements, formats, and languages. Companies spend billions of dollars on pharmacovigilance every year, reporting adverse drug reactions to the FDA in accordance with 7- or 15-day service level agreements.

Timeliness is a key aspect of the challenge. Companies must report adverse drug reactions to a country's regulator within a specified period, which might be as short as seven days for serious, life-threatening reactions. With millions of data points being collected, identifying these reactions can be time-consuming and difficult, and failure to do so can cost millions in fines. In fact, a recent study found pharmaceutical companies spent *\$4.87 billion* addressing pharmacovigilance in 2019, with the top 15 companies paying *\$1.2 billion* in fines for non-compliance in that year.

A study of 1.6 million adverse event reports found that more than 10% were not received by the FDA within the 15-day timeframe required by federal regulations, and more than 40,000 of these involved patient deaths. Three percent of reports were reported three to six months late, while another 3% were over six months late.

It is clear that pharmaceutical companies struggle with manually detecting, processing, and classifying adverse event cases. Manual processes take too long, especially with millions of data points streaming in daily from hospitals, legal channels, social media, and other sources. Even worse, the inflows arrive in different formats, complicating reporting to regulators. With case numbers increasing, an efficient, scalable, automated approach was urgently needed. Developers of the new approach faced a challenge on four fronts:

- **Pharmacovigilance spend:** Companies already spend billions of dollars on pharmacovigilance every year to report adverse drug reactions to the FDA within the agency's 7- and 15-day SLAs.
- **Saving lives:** Adverse drug reactions account for nearly 7% of hospital admissions in the U.S. alone, according to some estimates, with 50% of them being avoidable.
- **Costly non-compliance:** Millions of dollars are paid in fines and penalties by pharmaceutical companies due to delayed reporting of serious, fatal, or life-threatening drug reactions.
- **Big data challenge:** The industry continues to struggle with ingesting huge amounts of patient data and classifying cases as serious or non-serious within a very short period of time.

Solution

While existing pharmacovigilance spending presents a severe challenge, the application of AI and automation can speed up the extraction and processing of adverse events. Achieving this acceleration starts by targeting a key activity in the process: Ingestion. A data warehousing-led platform can accept and process the vast amount of data streaming from multiple sources. From there, companies can attack the challenge in stages. The key steps are as follows.

Data Governance

Start with a data governance model and a data ingestion framework, taking in data in many formats, such as a Word file, a PDF, a hospital image, or a doctor's handwritten note. These inputs are converted into a format readable by the platform.

Data Standardization (Extraction, Transformation, and Loading)

From there, the readable format runs against optical character recognition (OCR) and natural language processing (NLP) to extract keywords, creating a "universe" of keywords for use in these reports. The keywords are compared against a short list in a library and any matches are extracted to be investigated further.

Once data is ingested and keywords extracted, the platform must make sense of it all. Context is key: For example, a statement made in jest about a drug could easily be understood correctly by humans but misinterpreted by AI.

Data Splitting and Data Curation

It is reasonable to assume the incoming client data for training and validation is being processed and ultimately stored as emails and/or emails with attachments. If you were to start with 100 emails received over the span of a few years, you might end up with about 80 that arrived with attachments. After deduplication, that number shrinks to 50 emails with unique attachments. We now could split this curated data into three sets: 80% being employed as training data, 10% as test data, and 10% as validation data. That works out to 40 emails for training, five for testing, and five for validation. Of course, each organization is different and these ratios should be refined to the specific scenario.

Model Building

Data scientists, in the meantime, can adopt the following approach and techniques to create an ensemble of models for processing the incoming data. **Figure 1** shows an example of the relationship between techniques (shown on the circle) and activities.

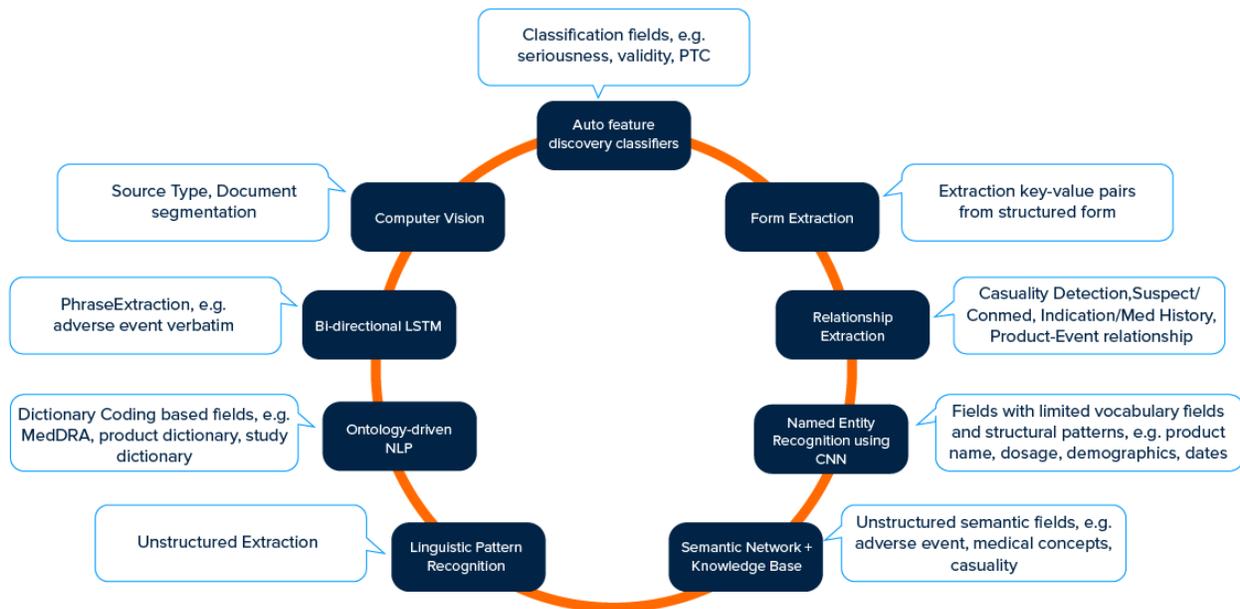


Figure 1. Proposed Models Ensemble—Achieving Predictions via Classification

As **Figure 1** shows, the following challenges and activities have been addressed by the techniques listed:

- Adverse event verbatim language: Bi-directional LSTM
- Source type or document segmentation: Computer vision
- Evaluating key value pairs: Forms extraction
- Causality detection: Relationship extraction
- Product name, dosage, and so on: Named entity recognition using convolutional neural network (CNN)
- Product dictionary-based coding: NLP

Once this step is complete, the resulting ensemble of models is trained to search for a set of associated keywords that can flag a report as serious or not serious. For pharmaceutical companies, the engine flags a serious report for attention based on its diagnosis and correlations, enabling them to report serious issues to regulators, such as the FDA, with no delay.

Training these models would require at least two years of historical data, with additional, continuous re-training of the models every six months based on newer data so that the models do not decay. These updates ensure that the models continue to provide reliable classifications of serious and non-serious events.

Keeping Company Models Distinct

If this solution were to serve multiple pharmaceutical clients, an obvious need arises to keep each company's proprietary data and model learnings siloed. Intellectual property issues and conflicts, for example, need to be actively managed in such an environment, with the learnings from one company's model fully dissociated from any others to avoid intellectual property claims, as well as to allay client concerns that their proprietary data could be exposed to third parties.

The concept of *federated learning* can address these legal concerns. Federated learning involves tapping into data sources from multiple companies, but *anonymizing* them by removing personally identifiable information (PII). Models are then trained on the results. The data is more broadly applicable as it comes from many different companies instead of just one. This also has the benefit of dramatically increasing the model's predictive powers and improving the reliability of its predictions.

The result is a "win-win" for all the companies as they benefit from more intelligent models. The learning is derived not just from one company's data, but from many. That benefit is shared by any company participating in a federated learning platform. If a bad reaction to a generic drug happens in a drug trial, it is likely to happen with other trials as well. This knowledge allows pharmaceutical companies to take preventative action.

Result

How significant is the impact of this technology? One estimate finds that by using this platform, an average-sized pharmaceutical company could achieve a potential reduction of FDA/regulatory penalties in the range of \$8 million to \$12 million per year.

There are operational cost savings as well. Organizations stand to save another \$5 million per year by increasing the efficiency of adverse event detection by 50% through the automation of case processing. Additional reductions in IT costs—around \$1 million per year—are also expected, based on reduced spending on cloud, infrastructure, and licensing in the range of 30%. In total, the platform could potentially save a company between \$14 million and \$19 million per annum.

Beyond direct cost savings, companies that leverage AI for pharmacovigilance can expect to see a number of other benefits. Among the impacts:

- **Speeding response:** Automatic extraction, coding, and processing of adverse event data results in quicker data processing and tighter reporting timelines.
- **Saving lives:** Rapid reporting protects patients by enabling timely interventions in the event of adverse drug reactions.
- **Improving compliance:** Removing delays allows pharmaceutical companies to comply with FDA mandates and preserve their reputation with regulators and the public at large.
- **Freeing resources:** Harnessing AI to transform pharmacovigilance into a sustainable and scalable

operating model frees up resources to focus on innovation.

- **Streamlining IT:** Robust automation results in reduced internal IT costs.

Lessons Learned

There were several lessons learned during the process of developing this solution.

Organizational Set-Up

Companies often park projects like this one under engineering or architectural leadership and steering, which can be a recipe for failure. There are plenty of organizational issues to consider, including:

- Tech teams often lack domain expertise and are removed from the business value of what they are building. Engineering teams typically have their own priorities and ideas on how to architect a solution (for example, favoring low-cost cloud resources and integration of model pipelines into different environments) while the data science team focuses tightly on building models that take input and produce a desired output. The data science folks will most likely not know how the output is consumed, because their job ends when models show a high level of accuracy.
- A product manager can define the features, elaborate on why one is more important to clients than others, and describe how machine learning model outputs are interpreted by the business. Yet at times, achieving high accuracy/F1 scores is not what's best for the business, and this is where the guidance to and coaching of the data science teams is critical.
- A product management-led approach can help join the engineering and data science teams. Crucial to this collaboration is ensuring that the teams communicate effectively to get the best results from the platform.
- Engaging UI/UX developers and designers up front can help the team nail the visualization, workflow design, and look and feel of the solution so the system supports users' efforts to navigate the analytics.

Ultimately, it's a good idea to establish a domain-expert business product manager who can guide the organization on what needs to be built and how it is critical to the success of the program.

Model Training

To achieve a reliable level of prediction, one needs to train the platform with millions of records involving both serious and non-serious effects. This training cycle is exceptionally long, so be prepared. One way to reduce time to market is to use an external training platform such as AWS SageMaker. This platform includes a component called Medical Comprehend, whose GPU processing power can produce results in a matter of days or weeks, instead of months or years.

Lifecycle Mismatch: AI Development vs. Scrum

Be aware that an out-of-the-box Scrum process *will not work* where AI development is involved. It is important to develop a modified software development lifecycle framework to address the fact that data science work is mostly experimental. It simply may not produce a final feature in every two-week sprint.

Product development often happens in sprints, while data scientists take months to create, train, and validate a model that needs to be integrated into the product. This mismatch creates timing issues around product releases and requires a product manager to define a unique development framework that allows for more flexibility beyond an ideal scrum methodology.

About Amit Arora



As a product innovation leader with 20+ years of global experience, Amit Arora has been at the forefront of new product creation and product-portfolio management at companies like Swiss Re, Genpact, Cisco, Virgin Mobile, Vodafone, and GE Capital, and has cumulatively generated more than \$500 million of new revenue by taking to market more than 30 new digital products. An early adopter of artificial intelligence (2010 onwards) and automation (2004 onwards), Arora is passionate about incorporating new technologies into his products, and has led technology build-vs-buy decisions for his employers. Arora is an affiliate member at the Data Science Institute at Columbia University and leverages this association to conduct research on use of AI and machine learning in new product development.

Arora obtained his MBA from Columbia Business School, attended Columbia Law, holds a Masters in Operations, an Engineering degree in Electronics, plus certifications in Agile and Six Sigma Practice. As a guest faculty at Columbia University, he regularly teaches data analytics and data science subjects to both graduate and undergraduate students. Amit has been invited regularly to speak at media platforms like The CBS News, The AI Summit, and O'Reilly's AI Conference, to name a few.

About GigaOm

GigaOm provides technical, operational, and business advice for IT's strategic digital enterprise and business initiatives. Enterprise business leaders, CIOs, and technology organizations partner with GigaOm for practical, actionable, strategic, and visionary advice for modernizing and transforming their business. GigaOm's advice empowers enterprises to successfully compete in an increasingly complicated business atmosphere that requires a solid understanding of constantly changing customer demands.

GigaOm works directly with enterprises both inside and outside of the IT organization to apply proven

research and methodologies designed to avoid pitfalls and roadblocks while balancing risk and innovation. Research methodologies include but are not limited to adoption and benchmarking surveys, use cases, interviews, ROI/TCO, market landscapes, strategic trends, and technical benchmarks. Our analysts possess 20+ years of experience advising a spectrum of clients from early adopters to mainstream enterprises.

GigaOm's perspective is that of the unbiased enterprise practitioner. Through this perspective, GigaOm connects with engaged and loyal subscribers on a deep and meaningful level.

Copyright

© [Knowingly, Inc.](#) 2021 "*AI-Driven Drug Safety*" is a trademark of [Knowingly, Inc.](#). For permission to reproduce this report, please contact sales@gigaom.com.