PharmaSUG 2023 - Paper SD-167

Bringing it All Together: Applying the Analytics Life Cycle for Natural Language Processing to Life Sciences

Sundaresh Sankaran[1], MS; Thomas W Sabo[1], MS;
1 SAS Institute Inc.

ABSTRACT

Life Sciences stakeholders, including the pharmaceutical industry and regulators, seek timely, accurate, and interpretable safety signal analysis. However, much of this data is unstructured freeform text, including physicians' notes and patient narratives. Manual review of such data is time-consuming and labor-intensive. It involves considerable effort spent to understand the context behind safety signals. Volume pressure and lack of standardization in review increases the chance of error, which can have expensive repercussions through regulatory holds, financial penalties, and reputational risk. Artificial Intelligence techniques such as Machine Learning and Natural Language Processing (NLP) help achieve better analysis. These capabilities are accessible to subject matter experts and programmers.

In this paper, we demonstrate the application of such techniques on SAS Viya including SAS Visual Text Analytics and show how they lead to faster realization of actionable insights. Automated dashboard explorations help domain experts and non-technical analytics consumers obtain early insights regarding topics and themes across freeform text, such as patterns of drug side effects or successfully treated symptoms. We further address the important issue of interpretability, to help organizations understand the logic behind models, and give subject matter experts the opportunity to refine these models without being a data scientist. Ultimately, we aim to give readers a better understanding of how the analytics lifecycle, particularly for NLP, can be applied to life sciences missions for a variety of user personas.

INTRODUCTION

The healthcare and life sciences sector seeks intelligent approaches to detect and understand safety signals from a variety of unstructured data sources. This desire also stems from a realization that unstructured data[1] is ubiquitous across the industry – found within doctor's notes, adverse event narratives, clinical trial summaries, case notations, literature reviews, patents and even drug labels.

Review of unstructured data is time-consuming and labor-intensive, with considerable effort spent to truly understand the context behind safety signals. Unstructured data, by its very nature, tends to be messy. It is scant in clear and concise labels, differs in attributes of length and style, and manual assessment tends to be subjective. This increases chances of an erroneous assessment, a rather costly consequence when it comes to assessing safety signals. Adverse events are also reported in high numbers. It is understood that 1.3 million emergency visits to a hospital occur due to adverse drug events[2]. Review of these adverse events - within provider, pharma, and regulatory organizations – is carried out by a team of skilled safety reviewers who have expertise in a given medical area. It is impractical to assume that the supply of such reviewers may scale adequately to accommodate review needs.

Volume pressure and lack of standardization increases the chance of error, which can have expensive repercussions through regulatory holds, financial penalties, and reputational risk. **Therefore, it's**

---

[1] Unstructured data and "text data" are used interchangeably across this paper, which is primarily focused on generating insights from text. However, note that unstructured data also refers to other non-conventional data sources such as audio recordings (such as those from support centers and interactions) and data from devices (x-rays, sonograms etc.). In the authors' opinion, the interpretation of the term "unstructured" is always a little nebulous, since each data artifact tends to have some semblance of a structure (just not the conventional tabular structure one expects). We prefer the word "unstructured" over "text" since the concepts and techniques described in this paper are equally applicable to other forms of data.
[2] Centers for Disease Control and Prevention, https://www.cdc.gov/medicationsafety/adult_adversedrugevents.html

important for stakeholder organizations to ensure timely, accurate, and interpretable safety signal analysis.

Analytics and Artificial Intelligence methods such as [Natural Language Processing] can help achieve this objective. Natural Language Processing, especially, has grown in popularity and has entered the general discourse, thanks to recent innovations (such as the advent of ChatGPT) in some specialized NLP tasks such as generating text. It remains however, a specialized area of study, which has led to the development and adoption of highly focused tools. The growth of multiple tools and technologies leads to the emergence of silos, reducing chances of attaining projects goals. Smart organizations appreciate holistic, unified approaches, which are easily accessible by many types of user personas. For ease of understanding, such approaches can be described using one term - the [Analytics Life Cycle.]

THE ANALYTICS LIFE CYCLE

It's a bird. No, it's a plane. No, it's <insert superhero of choice>

Similar sentiments can be expressed when talking about the analytics life cycle. People may different definitions of the term. However, if adopted and executed well, this can be a powerful approach to implement organization-wide analytics initiatives and yield huge benefits. Organizations, as they grow in complexity, tend to center their attention on groups of activities, such as data management, modeling, reporting/ BI etc. in isolation. Resources who carry out such activities also develop focused expertise in their designated area. When they focus on such activities in isolation, organizations may start to lose the big picture and plunge into inefficiencies and dogma.

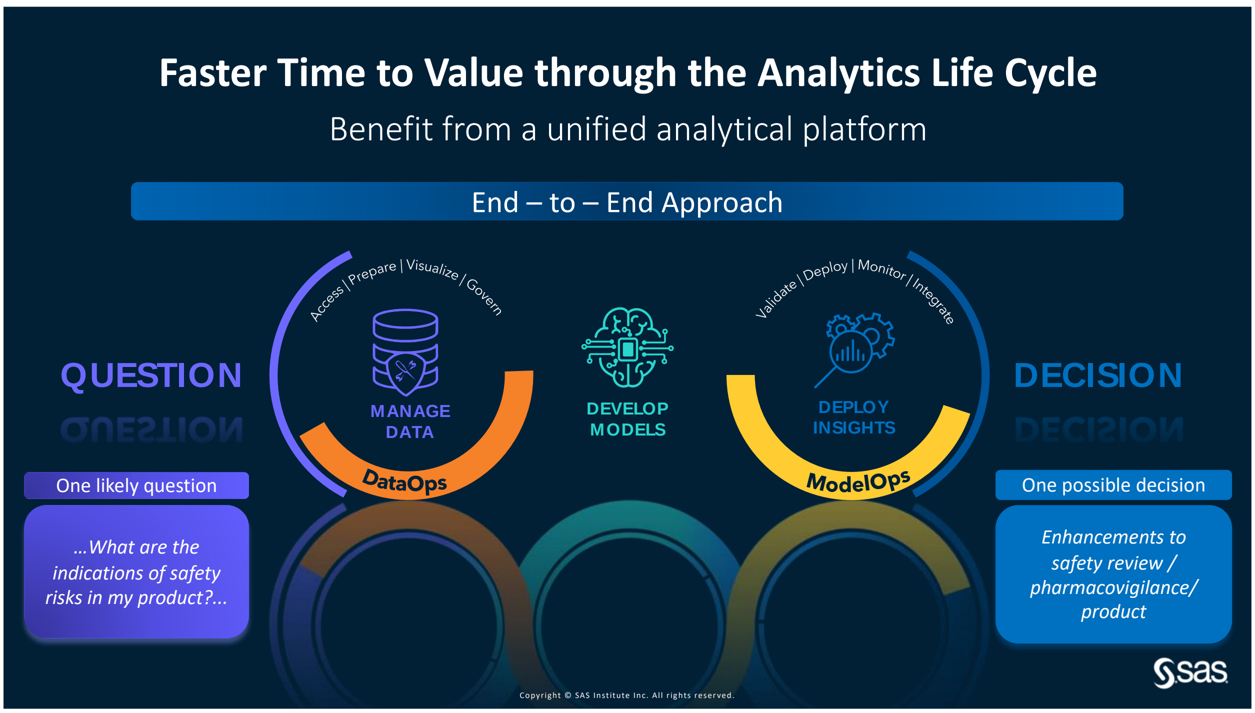

Figure 1 : Overview of the Analytics Life Cycle

As an approach, the analytics life cycle looks at analytical activities as a continuous loop with flow of information between three key phases of **data, discovery, and deployment**. A striking feature about the analytics life cycle is in terms of the context in which these phases are framed. Let's call them the two bookends of the analytics life cycle. These two bookends are **essential to provide focus and guide activities around data operations, AI, and model operations which comprise the analytics life cycle.**

1. The first bookend is a **business question.** It acts as the main motivation for the analytical activity to happen in the first place. A possible business question could be *"What are the indications of safety risk in my product / service?"* Healthcare & Life Sciences organizations grapple with this

question because of the very nature of their product / service, and the fact that any safety risks have huge downsides. The question could be reframed to be more specific across one product line (a new drug in the marketplace, or a new medical procedure or device) and, once answered, should spur further business questions downstream (such as *"How can I protect against the chances of adverse events occurring?"*)

2. The other bookend is an **insight-backed business decision.** This is the logical outcome which is designed after analytics has provided insights addressing the business question. In the context of this paper, it could be a signaling mechanism, such as a process which uses a Natural Language Processing (NLP) model to identify, classify, highlight & disambiguate safety risks, so that a reviewer could make better decisions with regards to reporting and mitigation.

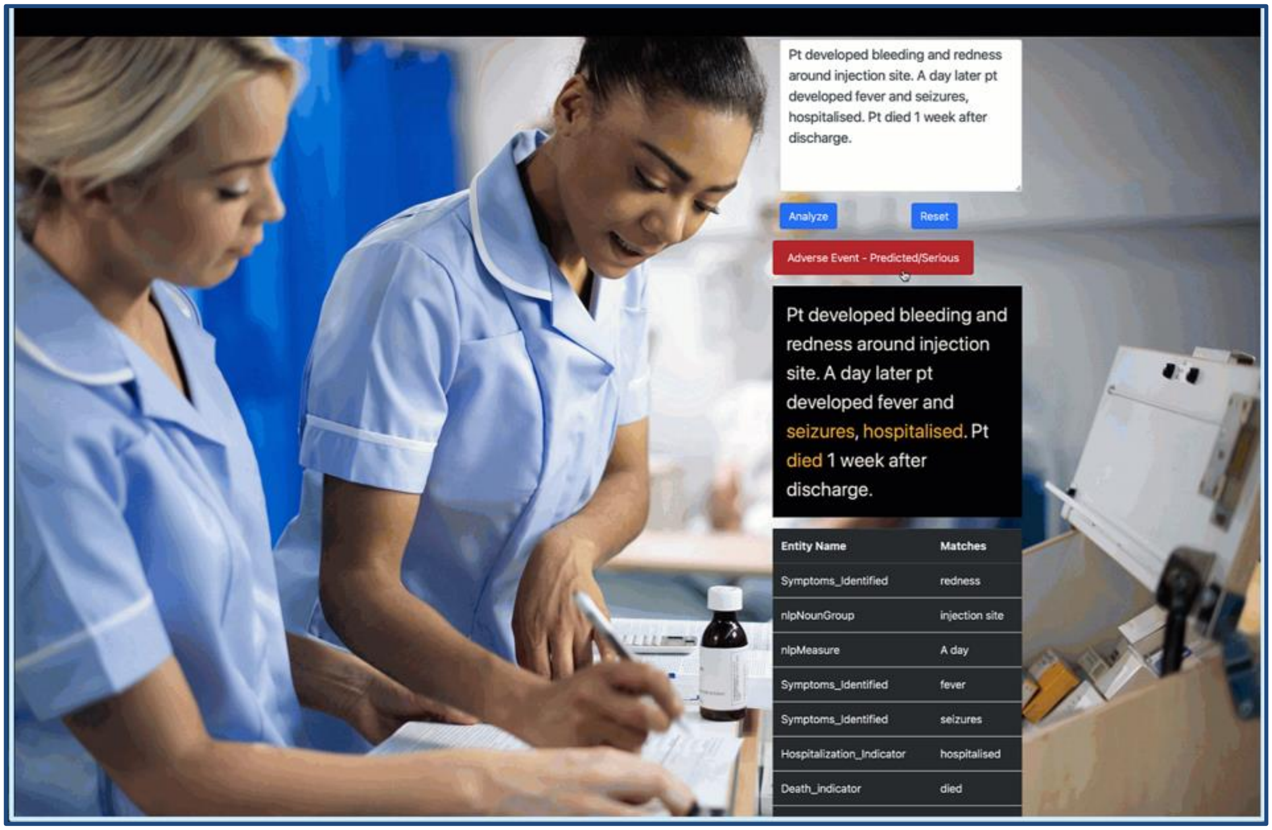

Figure 2 : An example of an operational decision based on an Adverse Events prediction model.

Let's look at such a decision outcome first, as depicted in the above image. This mimics a front-end application, which has processed a report which originated from a hospital and is now in front of a safety reviewer. Notice the prediction in the red box – Adverse Event – Predicted / Serious - which is instantaneously applied to the narrative indicating whether it has the potential to lead to a serious negative outcome or not[3]. Additional, contextual information, such as the entities which represent extracts from the report are also used by an AI model to suggest this signal and are used to provide a summarized depiction of the context around this case.

[3] This is a simplified example. Real-life adverse event narratives tend to contain longer passages with more details, such as references to drugs, progression indicators (day 1, dates etc.), medical abbreviations and many more.

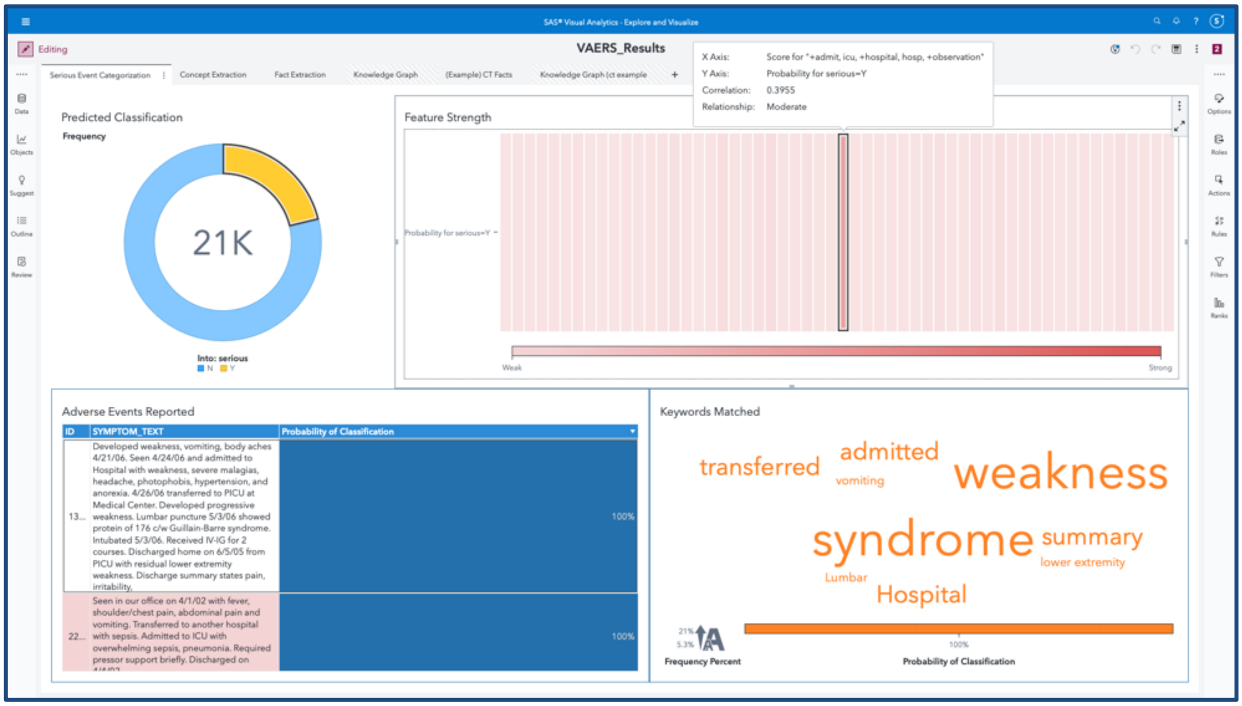

Figure 3 : A dashboard showing the results of predicting multiple adverse event narratives.

What you've noticed is only one example of a narrative going through a process known as triage. At the backend, what really happens is that many adverse event narratives get analyzed at scale, using a hybrid of machine learning techniques along with rules-based systems, all belonging to a class of AI called Natural Language Processing (NLP) or text analytics. NLP is extremely useful for text which (as a reminder) tends to be non-standardized, may carry subjective content, and differ in quality and structure.

Let's move away from processes and activities for a moment, and also consider the role of people, i.e., the resources who are part of an analytics & safety review organization. **Everybody should be able to participate in the Analytics Life Cycle**! Applications, methods, and techniques need to appeal to many user personas, including business users and programmers, not just specialized teams such as data scientists or data engineers. This has become increasingly critical due to two key factors: (a) a greater focus on explainability and transparency, and (b) gaps in availability of skilled analytical resources.

Greater adoption of machine learning-based models such as neural networks, gradient boosting etc., which operate in multiple iterations, has led to more accurate model predictions, but has increased complexity at the same time. At the same time, regulatory oversight of models and decision systems used in healthcare operations is quite intense. Business consumers of analytical models ask for more details around both how their models work, as well as more details around the process followed when creating these models. This enables them to answer technical and compliance-related questions in a transparent manner. Well-designed analytical life cycles should be able to help consumers of analytical models with such granular details.

Skills availability gaps make it tough to call upon expert AI practitioners to answer all questions. **Well-designed life cycles allow even business-oriented practitioners (such as a safety reviewer or a pharmacovigilance manager) to initiate analysis and gain rapid insights, with the flexibility of collaborating with experts on more complex tasks.**

AN ILLUSTRATIVE EXAMPLE

As an illustrative example, we depict major components of an Analytics Life Cycle using data available from the [Vaccine Adverse Events Reporting System (VAERS)]() website. VAERS reports contain both structured and narrative information about vaccine recipients, mostly young children inoculated against diseases such as measles, mumps, chickenpox etc., but also including adult vaccines such as the flu vaccine. To such reports, we have also added an additional column, serious which is a label indicating whether some events have led to serious negative outcomes (for example, the patient died, suffered loss of function, or was hospitalized for a long period of time) or not.

We chose SAS Viya® as the technology platform on which the analytical life cycle will be executed. SAS Viya is a modern, unified analytical platform from SAS Institute which offers capabilities in data management, machine learning, natural language processing, model governance, and ModelOps. A nice feature of SAS Viya is also its ability to offer multiple interfaces which cater to different user personas, a mix of Graphical User Interfaces (GUIs) and programming interfaces, and also access to both proprietary and open source (Python and R) algorithms from a common application.

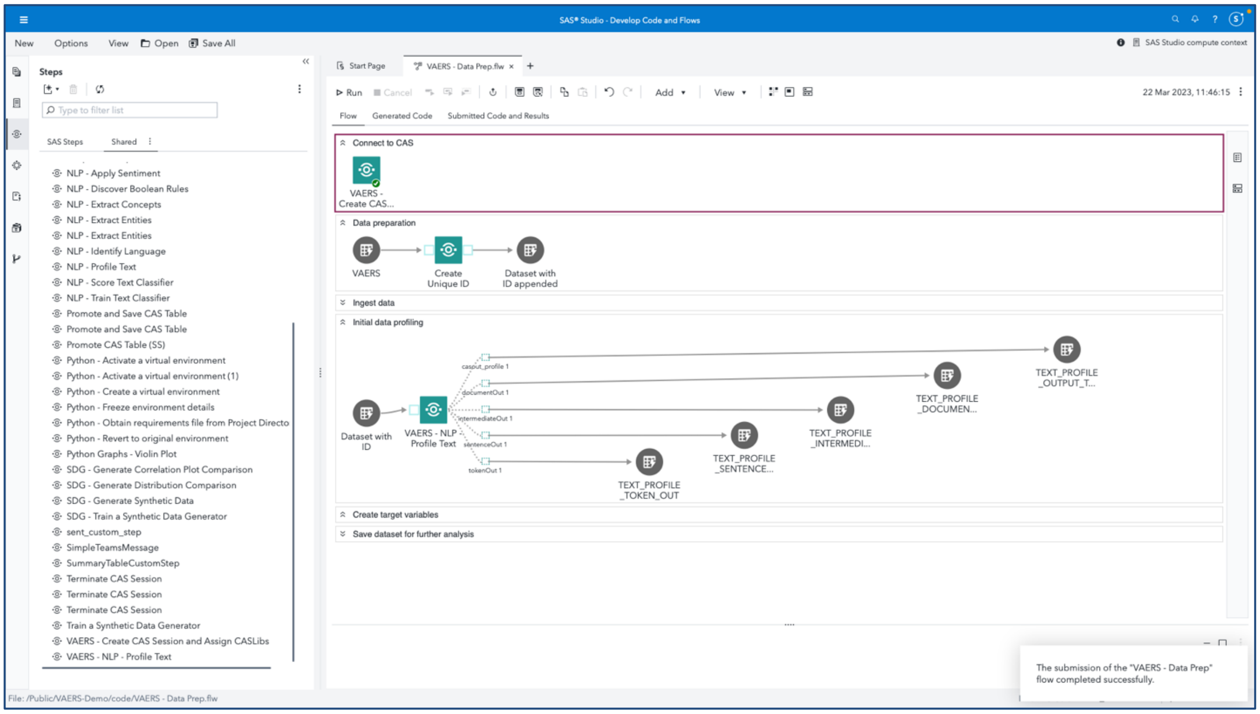

Figure 4 : A SAS Studio Flow, enabling data transformations involving low code components.

Let's begin upstream and start with data. We made use of SAS Studio Flows and Custom Steps, which are low code components that enable data engineers to easily ingest data from a variety of sources and transform them into a table which is suitable for analysis[4]. A variety of transformations are possible through this application, including creating new variables (such as a unique ID), recoding variables (creating a flag to identify serious adverse events), merging or appending additional datasets as required. In addition to predefined transformations available out of the box, many other custom steps are also provided as community contributions, such as a GitHub repository for SAS Studio Custom Steps (https://github.com/sassoftware/sas-studio-custom-steps#readme). These low-code components also enable a "shift left" of analytics – for example, some Natural Language Processing can be carried out early on in the analysis in order to provide an idea of the pre-processing that may be required on the adverse event narratives contained within VAERS[5].

Now that data is available in the system, notice that analysis can be initiated even by a safety reviewer who may not have deep technical knowledge of analytics or programming skills. Yet, they are also empowered within the analytics life cycle through automated components such as the Text topics object, which carries

---

[4] This is also referred to as an Analytical Base Table (ABT)
[5] For more details, refer "Profiling Text with SAS Viya" : https://communities.sas.com/t5/SAS-Communities-Library/Profiling-Text-with-SAS-Viya/ta-p/845109

out Topic Discovery. This can be found within the "Explore and Visualize" section of SAS Viya, within a self-service application called SAS Visual Analytics.

Topic discovery is an unsupervised technique, meaning, no prior judgements or hypotheses are imposed on the data. Rather, a statistical algorithm helps describe the topics or themes that are latent in the data – answering the question *"What **is** my data telling me???"*

The algorithm powering Topic Discovery is Singular Vector Decomposition (SVD). Singular Vector Decomposition applies a matrix factorization to term and document statistics from a corpus, in order to arrive at a representation of both terms and documents, as a vector of numbers. In short, given a narrative (containing X terms), and a specified number of dimensions (say, 300), the tmSVD action from SAS Viya outputs the following two tables.

1. A document embedding table, containing 300 features (columns) for each narrative.
2. A term embeddings table, containing 300 features (columns) for each term.

Table #1 from above is useful for carrying out downstream machine learning exercises to predict a serious adverse event. Table #2 also proves useful when it comes to harnessing the capabilities of deep learning[6], a subclass of machine learning which is more rigorous and a valuable addition to our arsenal of techniques. We'll interact with these features as we progress along the analytics lifecycle.

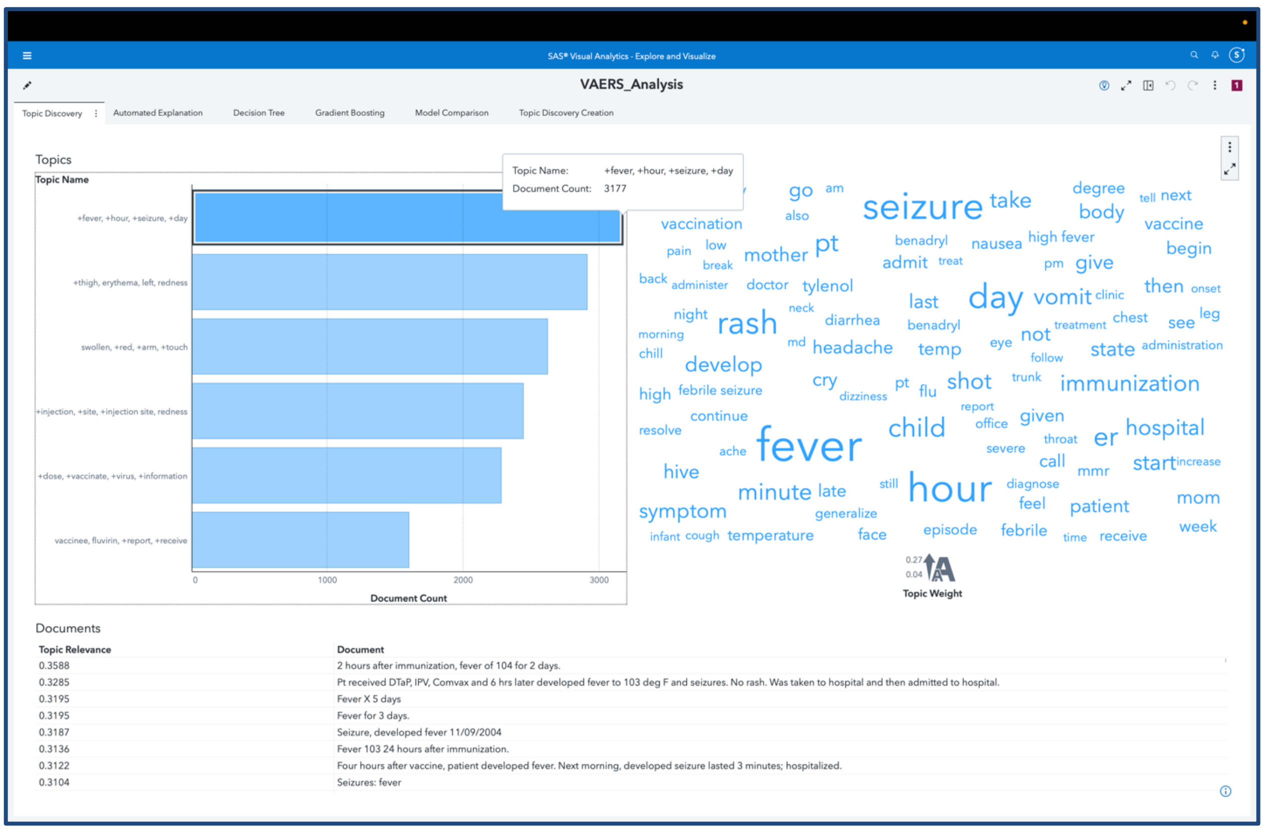

Figure 5 : Topic Discovery conducted as part of Exploratory Data Analysis

For now, the focus is more on facilitating ease of understanding and consuming results, which is carried out by the Topic Discovery report. For example, you note that around 3000 reports are characterized by a topic containing important terms such as *seizure, fever, hour*, etc., while another 2500 are characterized by

[6] Outlined in PharmaSUG 2023 Paper 110 (Deep Learning to Classify Adverse Events from Patient Narratives) in more detail.

a topic featuring terms such as *swelling, redness* etc. Terms and relevant statistics shown above are all output from the tmSVD action but are presented in a visual manner so that a non-technical user persona can obtain quick insights at the exploratory data analysis stage.

For a moment, let's close our eyes and regress to our childhood, where we were all dragged kicking and screaming to sit in front of a doctor who jabbed a needle deep into us. Traumatic as that experience may have been, it's highly likely that some of the swelling and redness some of us experienced may have been temporary at best. This is just a hypothesis, but a hypothesis that has been formed after exploration of the data has led us to understand that there are frequent references to "swelling and redness" type topics within these adverse event narratives.

We go further to prove this hypothesis. The previous analytical activity has allowed us to perform an important task – we have been able to convert text to a numerical proxy or proxies – a set of columns which indicate the relevance or embeddings for each narrative against the number of topics we discovered. Numbers are more amenable to analysis, as they are objective and can help us in constructing generalizable and scalable analytical models. Before doing so, we first seek to establish a relationship between these topics and the outcome variable, namely variable called "serious" which we had created during the Data Preparation stage.

The label **serious = Y** indicates any adverse event which has been tagged as either resulting in the patient getting hospitalized (with certain qualifiers, such as a stay in the ICU, long-term stay etc.), passing away, or suffering a loss of function or disability. **Serious = N** indicates all other narratives. Note that the choice of target variable depends on the business question you wish to answer (referring to the original bookends of the analytical life cycle) and therefore, within your control.

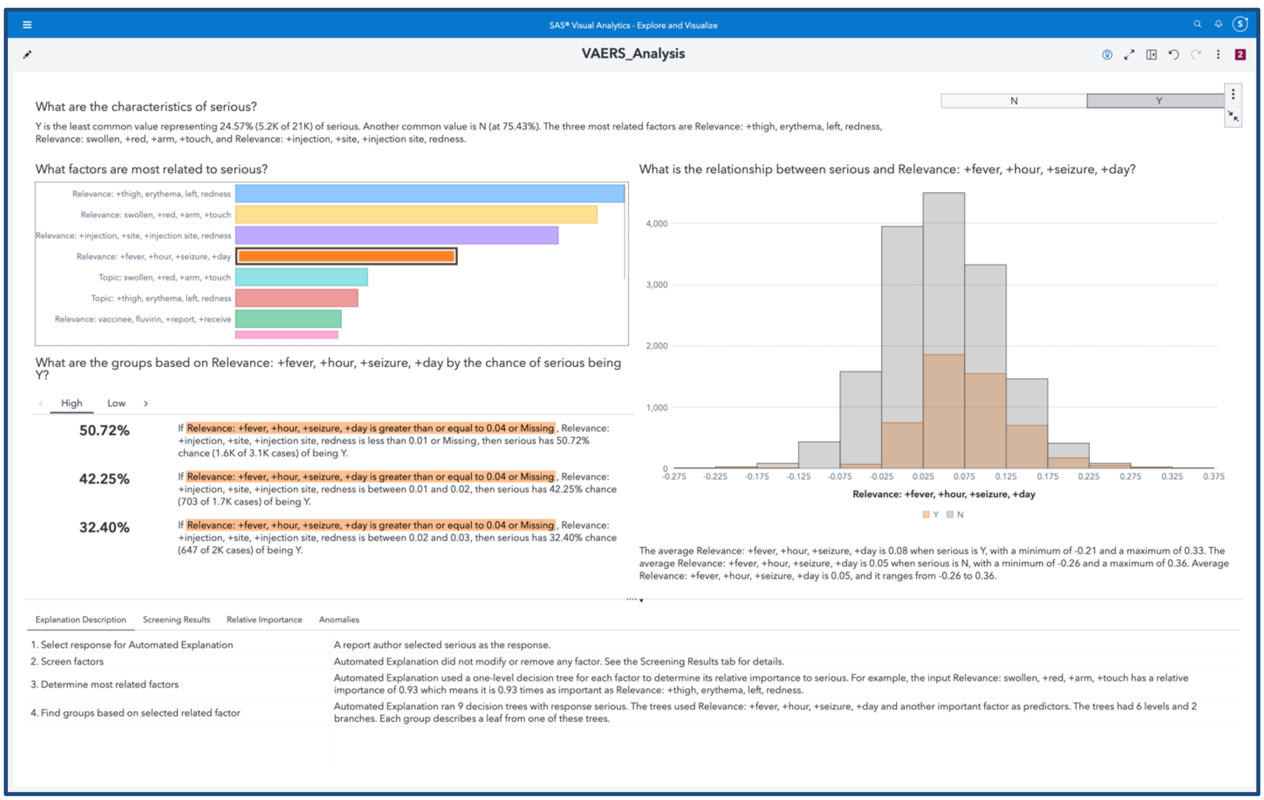

Figure 6 : Automated Explanation highlighting the "fever, siezure..." topic.

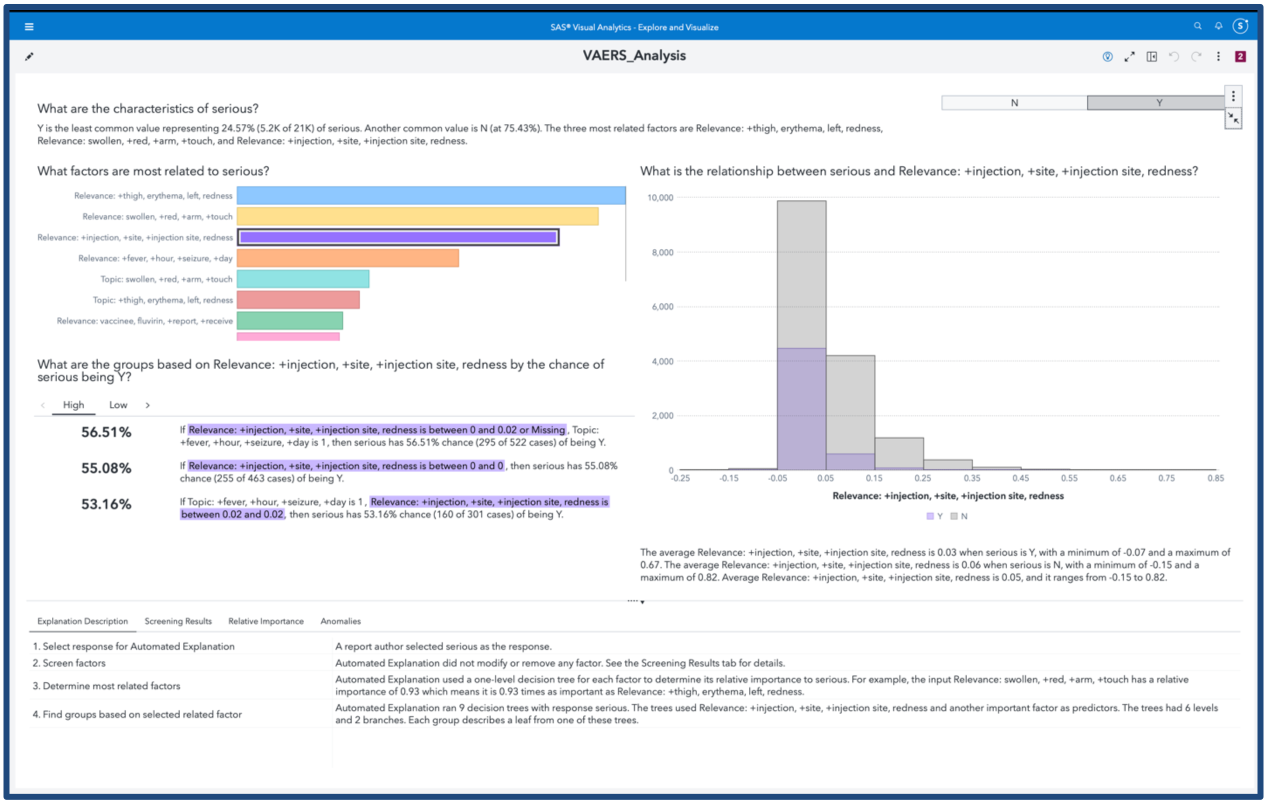

Figure 7 : Automated Explanation showing the "injection, site..." topic.

Through the charts above, you notice that the “seizure, fever...” topic is in fact directly correlated to serious AE outcomes, which is intuitive, while the “redness and swelling” topic seems to have an inelastic or even a negative relationship.

Unbeknownst to the safety reviewers, they have already entered the area of machine learning. **The Analytics Life Cycle facilitates the introduction of complex analytics techniques in a graduated manner, without sacrificing necessary complexity for simple messages.** The charts above (generated through the Visual Analytics Automated Explanation object) are created after running a decision tree algorithm against the text topic features which were derived during the Topic Discovery stage, with the **serious** label as the target.

Regarding accessibility, this entire analysis can be executed with zero programming knowledge or scripting knowledge, making it a great component within self-service analytical tools. **Good analytical life cycles should incorporate a healthy mix of such self-service tools along with more complex interfaces.** Also note the various helpful pieces of text which “tell the story”, using an automated explanation component. These are not for cosmetic reasons alone but play a vital role in furthering the course of investigation. An easily consumable insight such as the “Fever, seizure…” topic being associated with serious adverse events, helps the business user understand future areas of focus and a direction on what guidance to provide their collaborators – data scientists and data engineers.

The last point is important, because we should acknowledge that all analytics can’t be simple enough to be executed from within a self-service application. There might be areas where more complex analysis is required, which can be answered only through more rigorous and computationally intensive methods – areas that a data scientist is specialized in.

And onward we progress into the Analytical Life Cycle, deeper into the “Discovery” (the middle portion) phase!

We are now at a stage where the safety reviewer (remember, a person with greater domain knowledge than technical / analytical expertise) has been able to explore the data and gather some insights. It's also possible for them to carry out some basic classification activity utilizing machine learning models (such as decision trees, logistic regression, gradient boosting etc.). But they acknowledge that there's a limit to their exploration. The time has come to transfer information and control over to a team of data scientists, while making sure that everything need not be built from scratch at the same time.

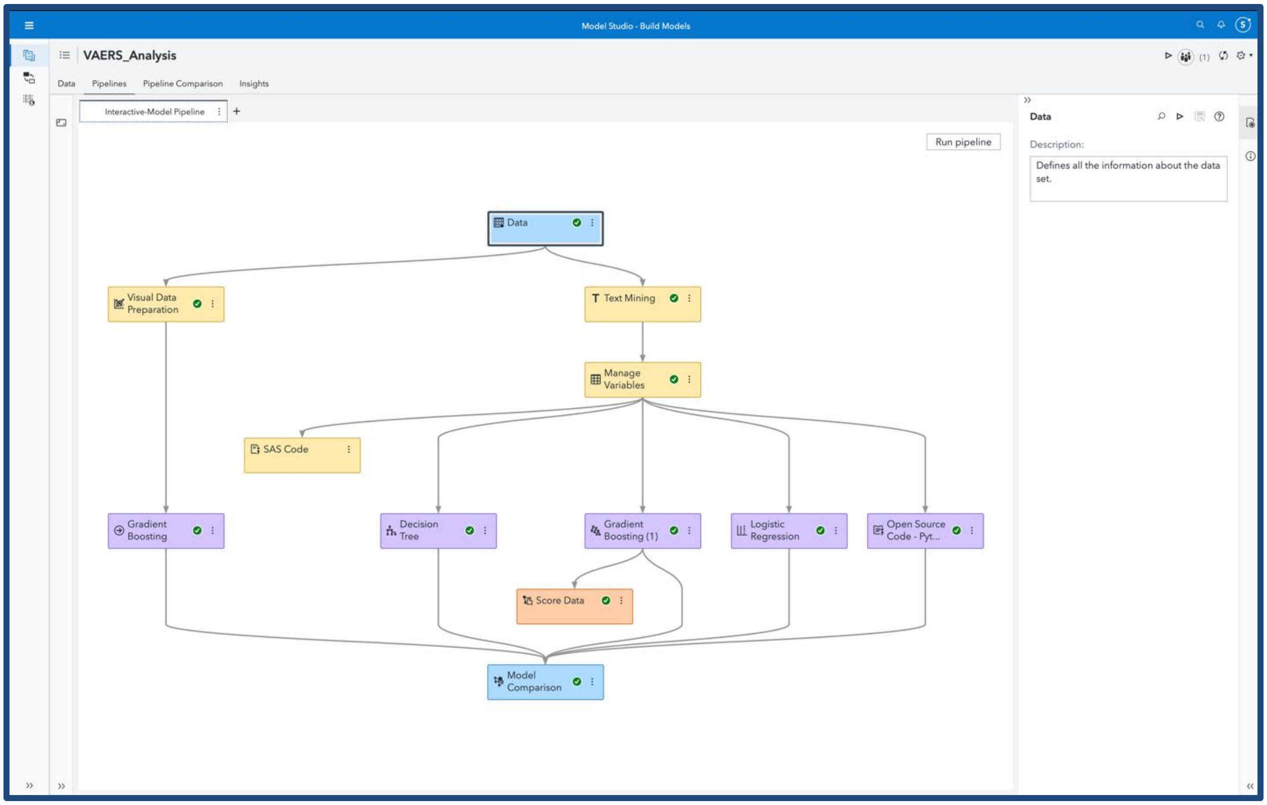

Figure 8 : A Model Studio project showing multiple models applied within a pipeline.

Well-designed analytics life cycles allow for seamless transfer of metadata and information within a collaborative workspace, which is a unified platform through which all applications are accessible. In this case, the platform is SAS Viya, and what you can witness in the above figure is that relevant information from the visualization application has been seamlessly transferred to the data scientist's workbench – the model development application. Seamless transfer of information ensures that the data scientist need not recreate all the steps of the business user – they just need to extend the existing analysis and flesh it out further. For example, text mining with 300 dimensions, which is much more granular and specific than the initial 6 – 8 topics discovered during exploration in SAS Visual Analytics.

Note that this entire pipeline of modeling algorithms is extensible. You can choose to add multiple proprietary algorithms to the pipeline to identify accurate models which identify serious adverse event outcomes with higher degrees of accuracy. In addition to proprietary algorithms, the analytics life cycle is also accommodative in nature and can also incorporate open-source algorithms from Python and R.

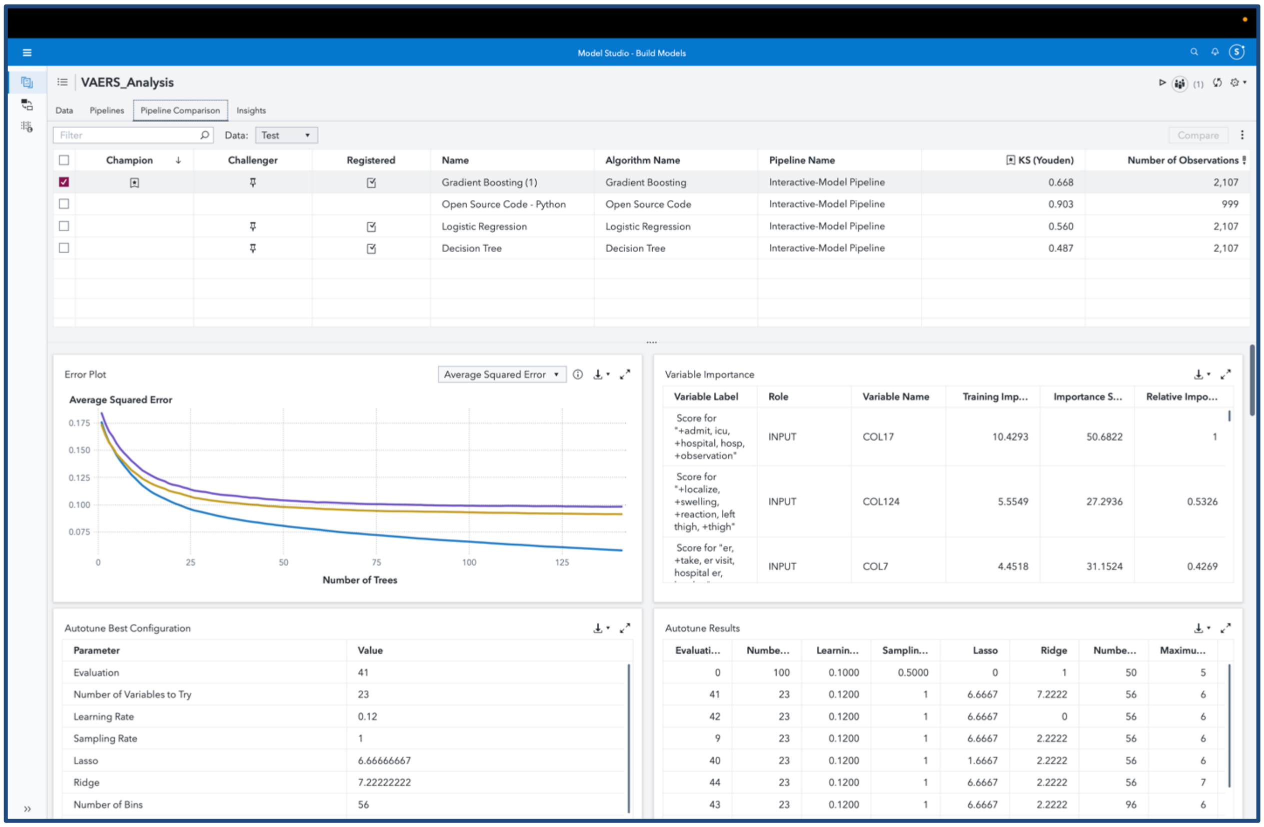

Figure 9 : Comparing results of multiple candidate models.

All these different models can be assessed using objective criteria through model comparison and pipeline comparison, which determines the best model (also known as a champion model). The model development interface should also allow for easy governance and shared templates, as is supported by Model Studio within SAS Viya, shown above. It's also possible to assign other models a status of Challenger, indicating that these also show good performance and can be considered for downstream deployment. Together, champion and challenger models are known as candidate models which are then registered in a model repository for governance purposes.

We make our way into the world of deployment. At this stage, candidate models exist, and have been assessed (from a technical perspective) by a model steward. When it comes to operationalization of the champion model, a major challenge often posed by safety reviewers to analytical developers can be stated as follows: *"..all this is great, but I am uneasy about using a model that seems very complex – it's like a black box to me! How can I understand what's happening?"*

Let's deal with the all-important topics of explainability and interpretability. Interpretability is a key demand of business organizations as analytics solutions increase in complexity. Nobody likes to just go ahead and sit in that shiny new self-driving car! Good analytics life cycles anticipate this challenge and provide heavy emphasis on automatically generated documentation which clearly explain the inner workings of the analytics model. Some examples of explainability measures include partial dependency plots, LIME & ICE plots, and HyperSHAP plots. In addition, interpretability of a model's effect should also be carried out keeping simplicity in mind. For this purpose, we take the predictions of these machine learning models, and run them by a rules-based NLP categorization model, in order to derive logical translations of the machine learning model's predictions in the form of Boolean logic. This is a further testament to the fact that **analytics life cycles accommodate multiple approaches**.

Instead of relying on a machine learning-based model alone, which does tend to resemble a black box in its inner workings, SAS Viya also allows us to use the predictions derived from the machine learning-model and run them through a rules-based text analytics model. This provides us a list of keywords which help explain the main drivers of why certain adverse event narratives are classified as “serious” versus “non-serious” – thus making the black box more of a ‘glass box’! This practice of combining multiple analytical techniques is known as “Composite AI” and helps promote the message that good analytics life cycles do not constrain the developer to a limited set of methods.

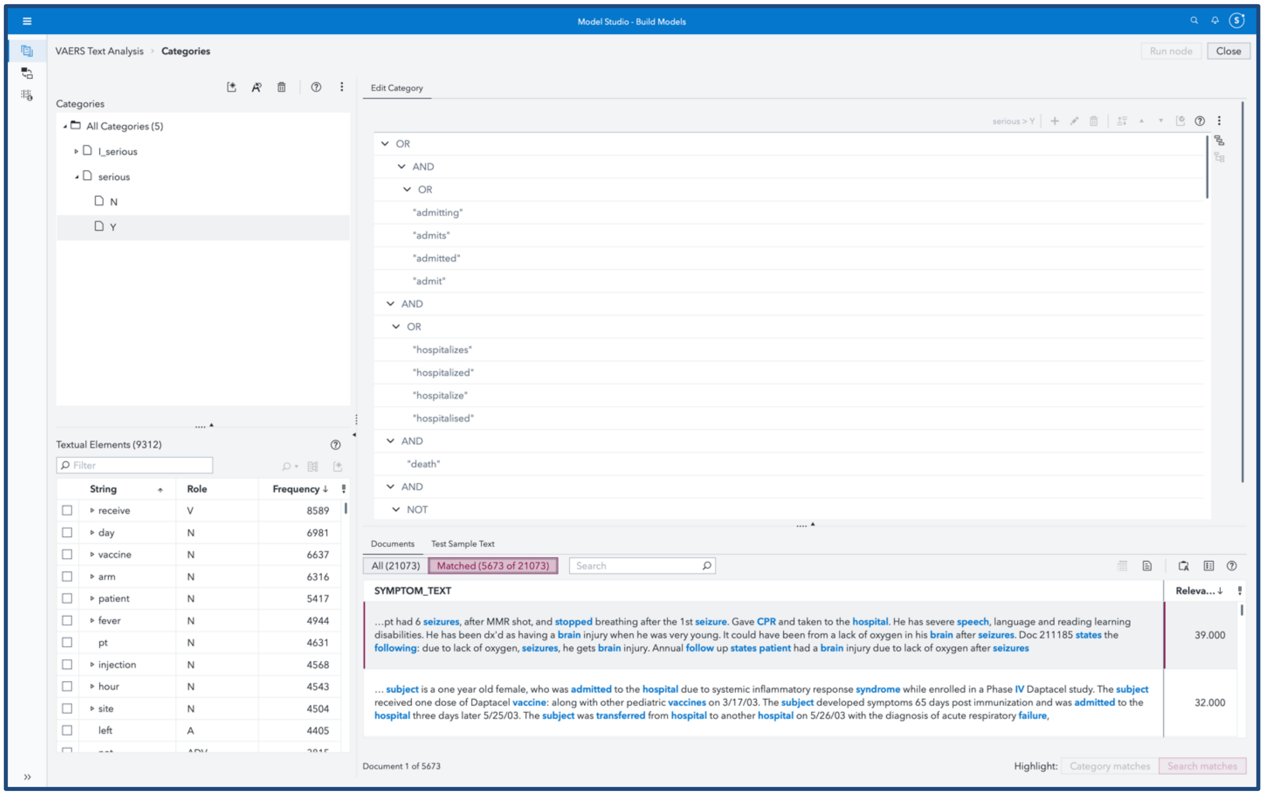

Figure 10 : Visual Text Analytics - Categories node

Two components at the end stages of the analytics life cycle now beckon to us. The first is the very act of operationalizing the model. Also known as scoring operations, inferencing, or publishing the model, this requires the final champion model to be made available in an environment where it can be scored through multiple modes. Popular mechanisms tend to be either “batch” scoring processes which run through a collection of narratives in one go, or “on-demand” scoring processes which may be run against one narrative at a time, either in a real-time or near-real time situation. Many organizations have communicated that they would like to reduce their dependence on IT organizations when it comes to model deployment, and good analytics life cycles should account for the same through allowing for a wide range of deployment mechanisms. In this example, the authors deployed the champion model as a module to a REST API service called Microanalytic Services, and this enabled the front-end application shown earlier.

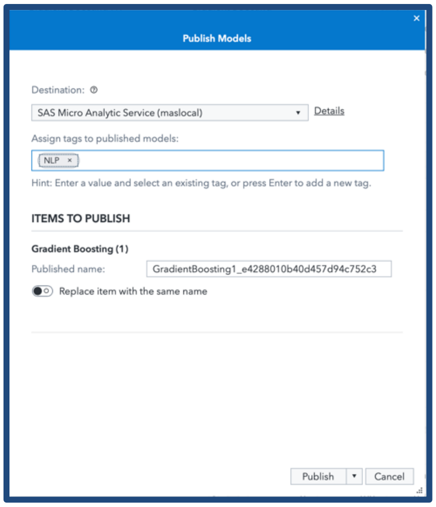

Figure 11 : Publishing a model to a REST API endpoint.

Finally, **analytics life cycles should also look to the future**. **Good analytics life cycles do not and should not assume that all models are good for perpetuity.** Models do deteriorate, due to a variety of reasons such as loss of relevance, a change in the underlying factors, portfolio shift, or shocks. It is important to identify such signs of deterioration at early stages, so that remediation actions such as a retraining of the model, or a recalibration of the weights could take place. Good analytics life cycles, like the image below, should incorporate reports that continuously monitor and alert model stewards regarding early signs of deterioration.

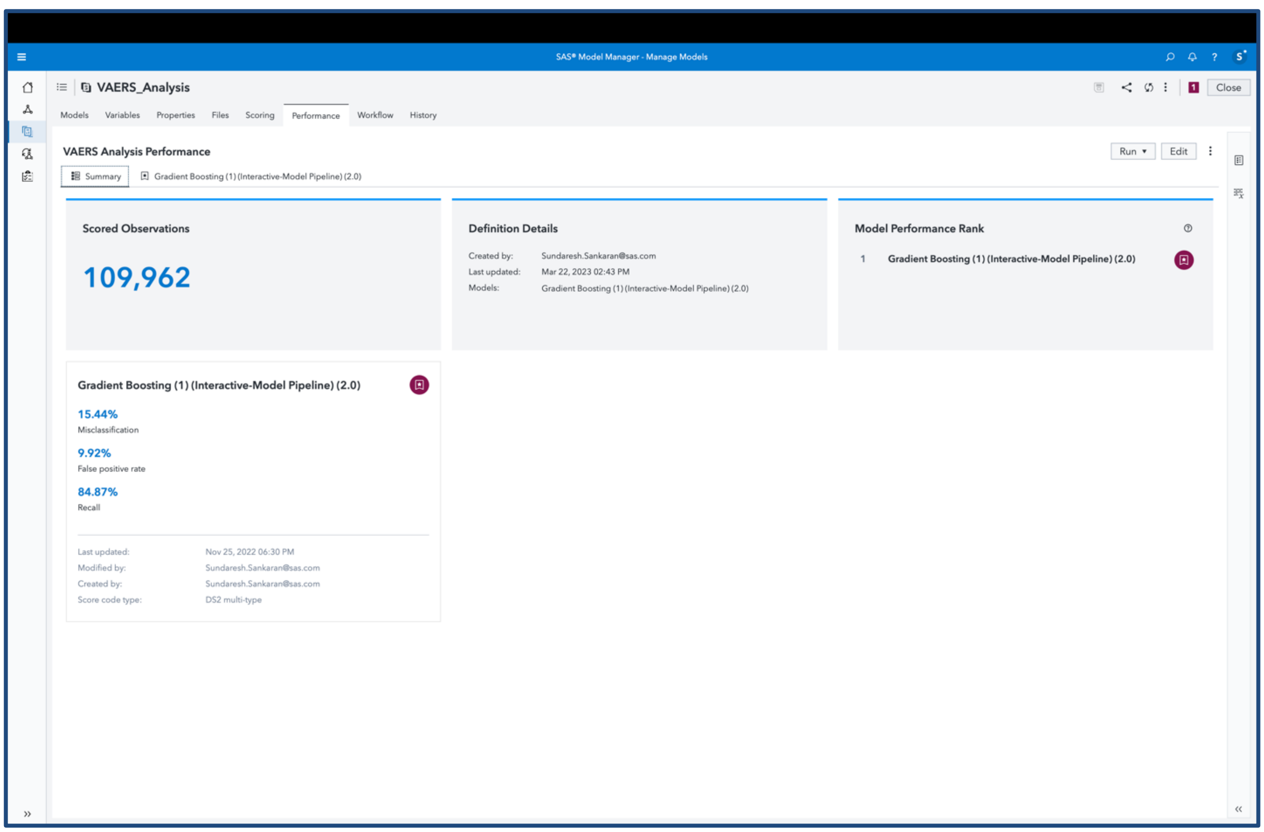

Figure 12 : A model performance monitoring report

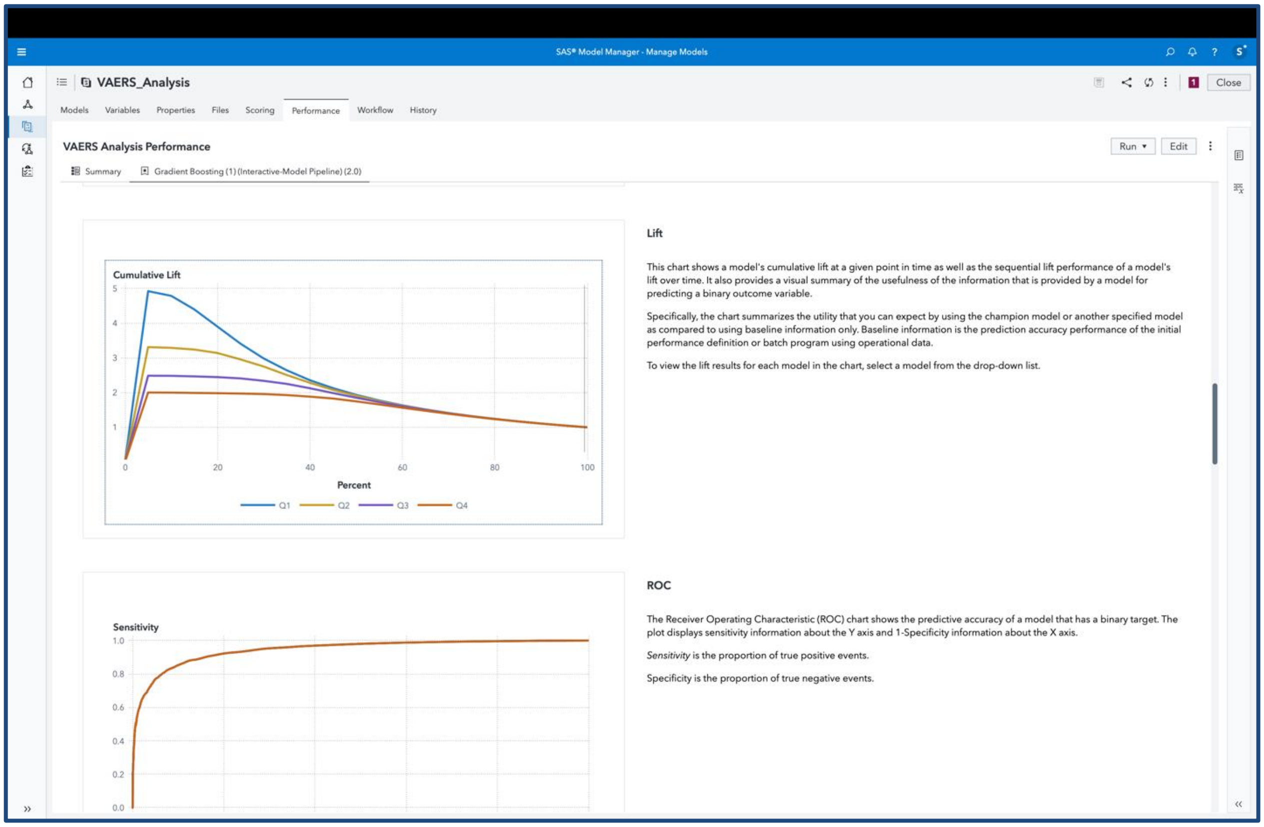

Figure 13 : Model performance monitoring report showing lift over quarters.

In the example represented from the figures above, we (hypothetically) deployed the VAERS model originally trained, and then tracked the same, using new data, quarter by quarter. This enables us to assess performance of the model (as represented by statistics such as misclassification rate, false positive rate, etc. over a long period of time, and make a determination whether the model still provides us benefits which are justified. Furthermore, we can notice a steady decrease quarter over quarter in terms of the efficacy of the model (through the Cumulative Lift measure), which indicates deterioration of the model, and is a signal for the organization to invest efforts in terms of improving the model (even to the extent of retraining a new model, if needed).

CONCLUSION

We have explored how the analytics life cycle is beneficial due to its **flexibility, suitability for multiple user personas, and accommodative nature**. We have specifically also explored the applications of the Analytics Life Cycle to unstructured data, in this case, applying Natural Language Processing to adverse event narratives. This has proven useful when considering the problem of extracting useful safety signals from unstructured text and provides downstream users rapid time to value. This approach is also portable and can be used, with minimal modifications, for other useful data sources in the life sciences domain such as medical literature, customer complaints and others.

We have also derived some useful principles about the right way to apply the Analytics Life Cycle approach, which we present below. As mentioned earlier in this paper, the Analytics Life Cycle could be defined in many different ways depending on interpretation. We hope that the principles we have drawn from successfully following the analytics life cycle will stand you in good stead.

1. The Analytics Life Cycle should be viewed as a **continuous loop between data, discovery and deployment.** Addressing these phases in silos will lead to inefficient outcomes.
2. Bookend your analytics processes with a business question on one end, and an insight-based decision on the other. **This will help provide focus and guide activities around data operations, AI, and model operations which comprise the analytics life cycle.**
3. **Everybody should be able to participate in the Analytics Life Cycle**! We want all user personas – business-oriented users, safety reviewers, data scientists, and data engineers – to be able to access the application in a unified and well-governed manner, in order to truly collaborate and jointly tackle a business problem.
4. **The analytics life cycles should be flexible -** Well-designed life cycles allow even business-oriented practitioners (such as a safety reviewer or a pharmacovigilance manager) to initiate analysis and gain rapid insights, with the flexibility of collaborating with experts on more complex tasks.
5. **A unified, end-to-end platform facilitates the smooth flow of the Analytics Life Cycle.** This makes it possible for the Analytics Life Cycle to facilitate the introduction of complex analytics techniques in a graduated manner.
6. **Good analytical life cycles should incorporate a healthy mix of self-service (visualization and exploration) tools along with more complex interfaces.** It should be possible to execute some AI techniques even from within self-service, low-code components.
7. **Well-designed analytics life cycles allow for seamless transfer of metadata and information within a collaborative workspace**, which is a unified platform through which all applications are accessible.
8. **Analytics Life Cycles accommodate multiple approaches**. They serve the cause of Composite AI, which recognizes the usefulness of multiple methods and approaches to solve a problem, while retaining control in terms of being able to compare and objectively choose between approaches.
9. **Analytics Life Cycles should consider the future**. **Good analytics life cycles do not and should not assume that all models are good for perpetuity.** Models can, and do degrade, and performance monitoring should be able to alert model stewards toward signs of deterioration at an early stage so as to enable proactive mitigation.
10. Every list of 9 points needs a #10, so we end this on a light-hearted note stating that Analytics Life Cycles, beyond a point, are just a convenient nomenclature to a set of practices which may need to adapt based on your organization's unique business problems. The authors try their best not to take the term too seriously and prescribe tasks and activities in a certain order (beyond what's logical) but attribute more importance to what the underlying unified platform should enable.

REFERENCES

1. Vaccine Adverse Events Reporting System (VAERS), co-sponsored by the Centers for Disease Control and Prevention (CDC), and the Food and Drug Administration (FDA), agencies of the U.S. Department of Health and Human Services (HHS).
2. Useful information about the Analytics Life Cycle : https://www.sas.com/en/whitepapers/manage-analytical-life-cycle-continuous-innovation-106179.html
3. Further Information on the SAS Studio Custom Steps GitHub repository : https://communities.sas.com/t5/SAS-Communities-Library/Custom-Steps-in-SAS-Studio-GitHub-repository-now-available/ta-p/840504
4. Sankaran, Sundaresh, 2022. Profiling Text with SAS Viya : https://communities.sas.com/t5/SAS-Communities-Library/Profiling-Text-with-SAS-Viya/ta-p/845109
5. Hatim, Sankaran, Sabo, McRae, Laufe, and Blanchard, 2023. Deep Learning to Classify Adverse Events from Patient Narratives, PharmaSUG 2023 Annual Conference, Paper 110.

6. SAS Help Center, 2023. What does the tmSVD action do? https://documentation.sas.com/doc/en/pgmsascdc/v_037/casactml/casactml_textmining_details06.htm
7. SAS e-Book for Visual Analytics. Chapter 10 – Advanced Topics – Performing Text Analytics : https://www.oreilly.com/library/view/sas-visual-analytics/9781952365119/Chapter_10.xhtml
8. Styll, Rick, 2020, Smarter and Faster with SAS® Visual Analytics. SAS Global Forum 2020 proceedings. https://support.sas.com/resources/papers/proceedings20/4609-2020.pdf
9. SAS Help Center, 2023. Concepts : Performance Monitoring. https://documentation.sas.com/?cdcId=mdlmgrcdc&cdcVersion=default&docsetId=mdlmgrug&docsetTarget=p1c6xm7tthdajkn1t6esm4n3kwnq.htm

CONTACT INFORMATION

Feel free to explore the analytics life cycle in more detail through the links provided throughout this paper, and reach out with questions.

Sundaresh Sankaran, Senior Solutions Architect
SAS Institute Inc.
O: (919)531-4921
C:(919)400-3266
sundaresh.sankaran@sas.com
https://www.linkedin.com/in/sundareshsankaran/

Tom Sabo, Advisory Solutions Architect
SAS Institute Inc.
+1 (703) 310-5717
tom.sabo@sas.com
https://www.linkedin.com/in/tomsabo/