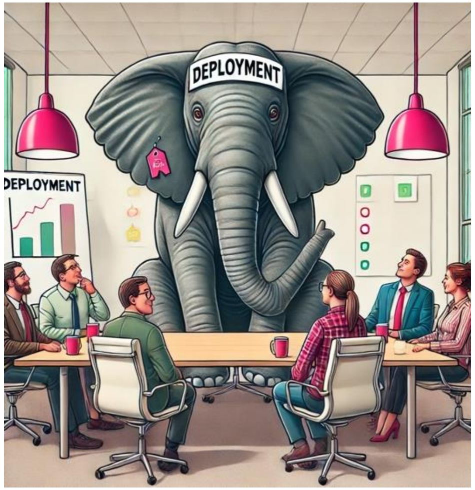

Is deployment the elephant in the room?

John McConnell

February 2026

FAQ's

- Is this session being recorded? Yes
- Can I get a copy of the slides? Yes, we'll email links to download materials after the session has ended.
- Can we arrange a re-run for colleagues? Yes, just ask us.
- How can I ask questions? All lines are muted so please use the chat panel
  - If we run out of time we will follow up with you.

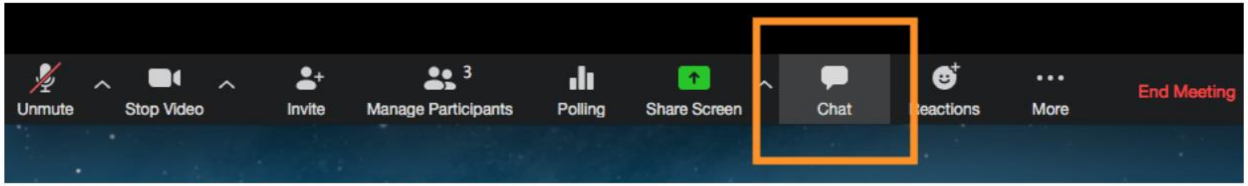

Is there a problem?

- **WHAT OTHERS ARE SAYING**

  "43% of Data Scientists say that 80%, or more, of their built models fail to deploy"
  - Eric Siegel – KD Nuggets January 2024

  According to Gartner analyst Nick Heudecker, over 85% of data science projects fail.
  - Ryohei Fujimak, DataNami 2020

- S&P Global Market Intelligence / 451 Research (Voice of the Enterprise), Oct 2025 Finds a big rise in "doesn't reach production": companies abandoning most AI initiatives before production increased from 17% to 42%, and the average organization scrapped 46% of proof-of-concept projects prior to production.

- **OUR OWN EXPERIENCE**

  We don't always succeed but have learned from our experiences

  We see more use of platforms and pipelines and more joined-up teamwork across functions

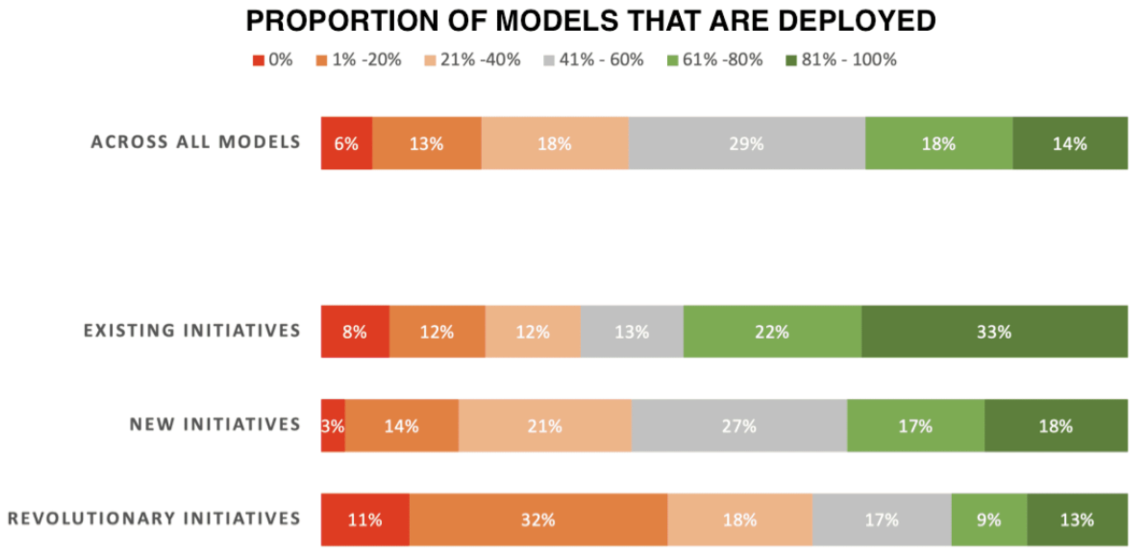

Key:

- **Existing initiatives:** Models developed to update/refresh an existing model that's already been successfully deployed
- **New initiatives:** Models developed to enhance an existing process for which no model was already deployed
- **Revolutionary initiatives:** Models developed to enable a new process or capability

Source: [Rexer Analytics] Data Science Survey 2024

The CRISP DM/DS process model

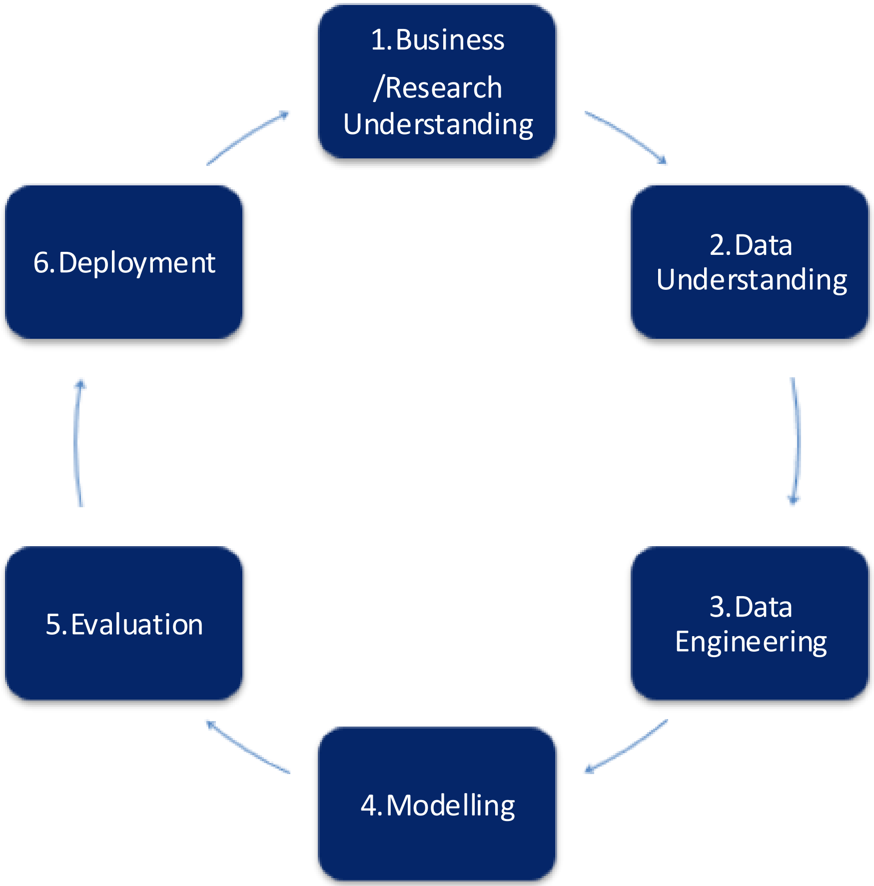

The Cross Industry Standard Process for Data Mining and Data Science.

For New Initiatives we look at the whole cycle

At the outset **Deployment** might seem somewhat far off … something we can think about if/when we have arrived at a successful model

1. Business/Research Understanding

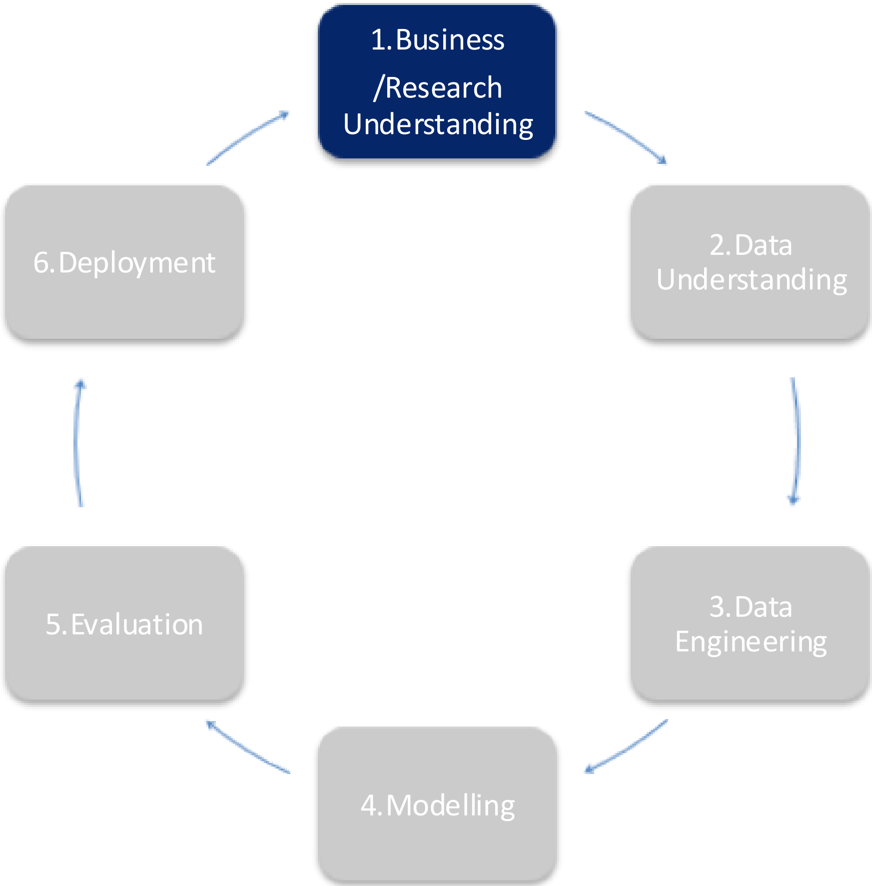

- We start with business/research conversations.
- Agree the objectives
- Understand the context
- Set success criteria
- Be sure to talk about "Deployment"

1. Agreeing Success Criteria

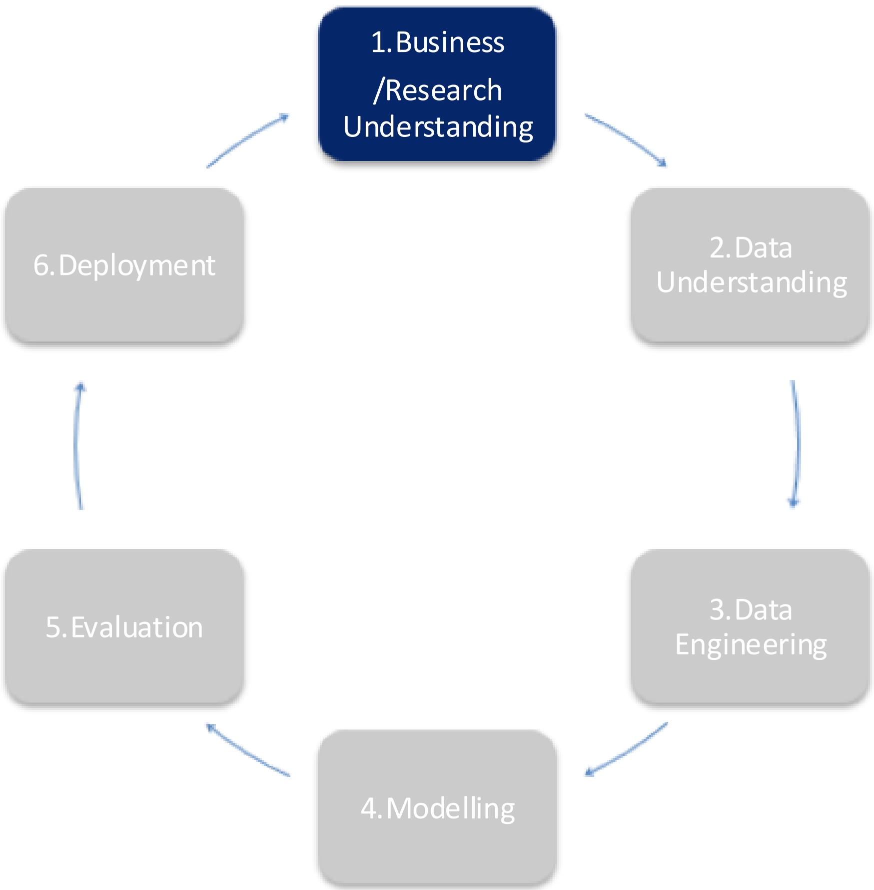

- An important step in the project where all stakeholders should align.
- For example, in a customer retention scenario, let’s say we can only operationally intervene with 2,000 customer at a time
- So we might agree that we want our model to accurately identify 70% of customers who are likely to leave within the top 2,000 most “at risk” (of leaving)
  - If our churn rate is 20% that is likely to be a strong model

6. Deployment

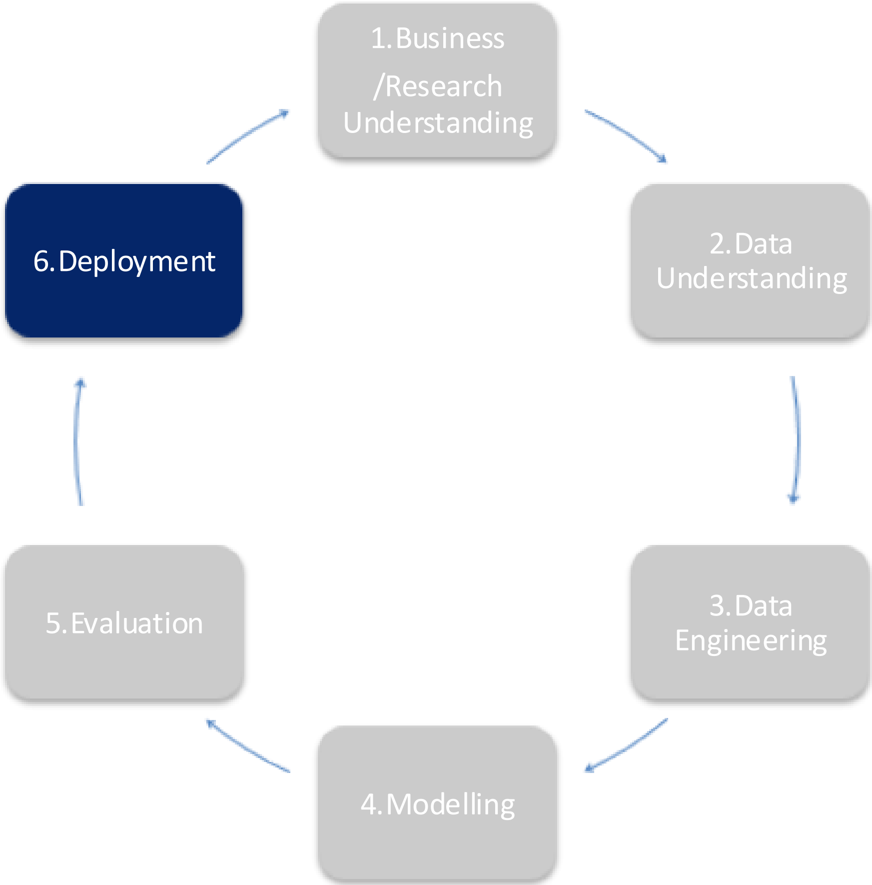

- The whole point of the exercise
- Can be about deploying, often strategic, insights that come from simpler (often statistical) models and statistical analyses
- More often about deploying models (and often the data engineering that feeds them)

6. Deployment

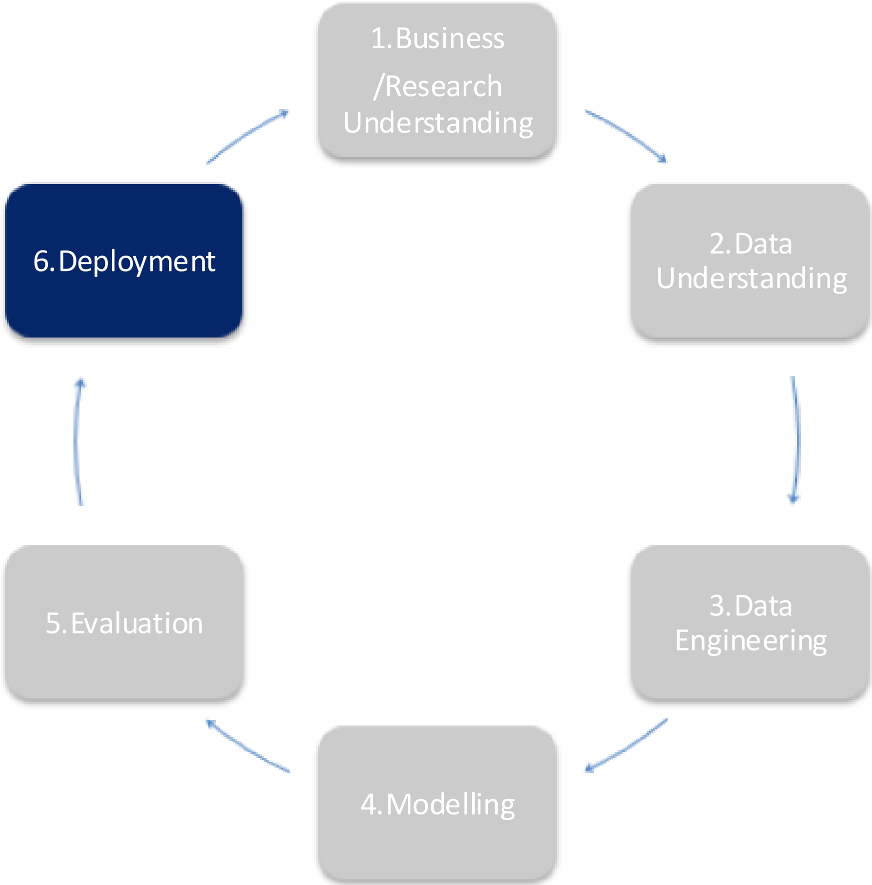

- Increasingly needs to be Real-Time
- Frequently needs to be scheduled
  - Typically processing batches
- Often integrated into operational systems
- Can be behind apps/platforms including
  - What-if? simulators
  - Decision Management tools

People (/ Roles)

Business

- Subject matter
- Decision makers
- Users

Data science

- Methodologies
- Analysis

Data

- Engineering
- Structure
- Subject matter

Technology

- Integration
- Building apps
- Platforms

BLOCKERS TO DEPLOYMENT SUCCESS

The Data

- We don’t have enough (when we assess in Data Understanding)
- The quality is too low
- It is too difficult to access, clean or prepare
  - The “effort to insight” ratio is too high
- What was accessible to us when we built our model is not available when we deploy

The Model(s)

- The model does not meet our accuracy target

OR

- It is essentially flawed
  - This could be data related, or issues in the data prep
- We may not realise it is flawed until it is deployed and it does not deliver the results we expected

Stakeholder alignment (people, politics)

- May block us during the Evaluation
- Often happens as/after we deploy
- Decision making stakeholders may not "buy" the model
  - Especially when we start a new initiative they often want to how the model is working
- **End user stakeholders** may not know how to use the model
- They may not want the model
  - It can be seen as a threat
  - The use of the model may be perceived to be adding to their workload

Technical alignment

- This most often hits right at the point of deployment
- Simply put, we can't move our model into production for a technical reason e.g.
  - The data sources we need for the model are not reachable or reachable in time
  - The target platform cannot run the model(s) (or we don't have one)
  - IT governance e.g. concerns about open source

(TECHNICAL) DEPLOYMENT OPTIONS

Option 1 – Recode the model into something else

Simple models can be (hand) coded

```
Call:
lm(formula = CONSUME ~ PRICE + INC + TEMP + PRICEINCi, data = datavar)

Residuals:
       Min         1Q     Median         3Q        Max
-0.0575279 -0.0163589 -0.0008483  0.0168662  0.0718922

Coefficients:
              Estimate Std. Error t value Pr(>|t|)
(Intercept)  0.1570203  0.2324673   0.675   0.5058
PRICE       -0.1636906  0.7438870  -0.220   0.8277
INC          0.0012301  0.0012133   1.014   0.3208
TEMP         0.0028231  0.0004171   6.769 5.31e-07 ***
PRICEINCi   -0.2786003  0.1344397  -2.072   0.0491 *
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 0.03094 on 24 degrees of freedom
Multiple R-squared: 0.7411, Adjusted R-squared: 0.698
F-statistic: 17.18 on 4 and 24 DF,  p-value: 8.968e-07
```

- Linear Regression models are a good example
- They fit easily into Excel formulae or into other code e.g. we’ve sometimes deployed models in SAP this way

Consume = 0.1570203 -0.1636906*PRICE + 0.0012301*INC + 0.0028231*TEMP -0.278600*PRICEINCi

Option 2 – Export the model into a code format that can be plugged into other software

Exporting SQL is one option (sometimes)

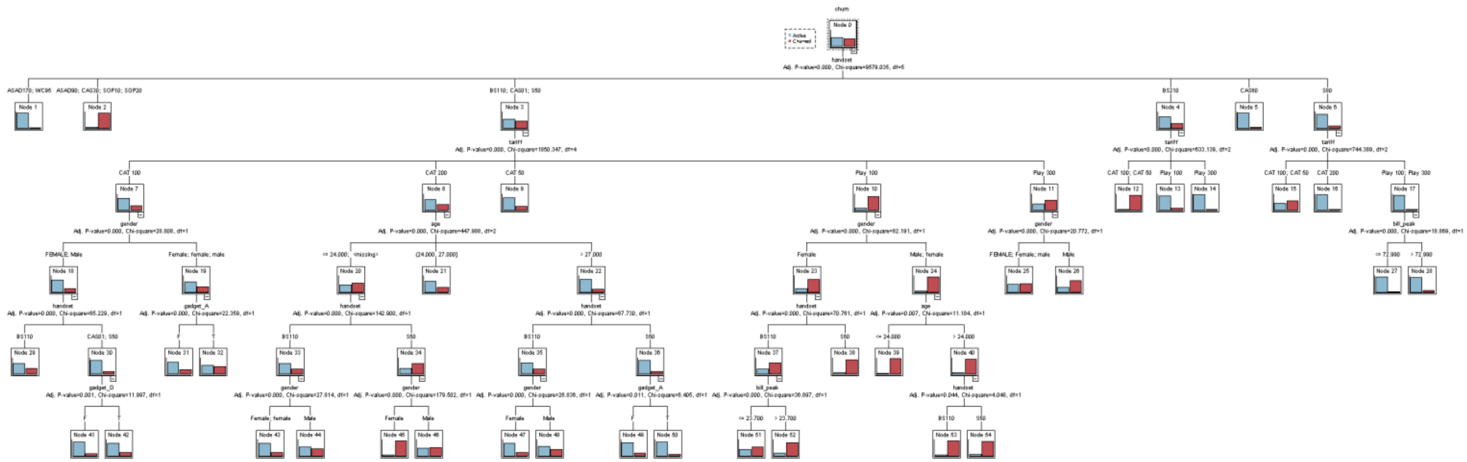

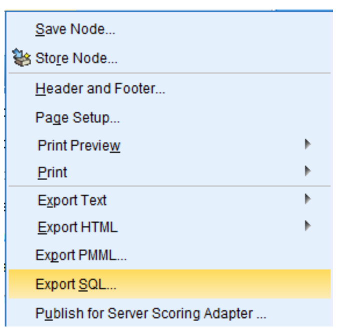

- Some tools support this. E.g. Python and R support it for some model types
- Caution around the **differences in SQL** between databases e.g. Oracle and IBM/DB2
- Some tools provide **target-specific SQL** which can sometimes be published into the database

PMML is the standard export/import format

www.dmg.org

Predictive
Model
Markup
Language

PMML 4.4.1 - General Structure

PMML uses XML to represent mining models. The structure of the models is described by an XML Schema. One or more mining models can be contained in a PMML document. A PMML document is an XML document with a root element of type PMML. The general structure of a PMML document is:

```
<?xml version="1.0"?>
<PMML version="4.4"
  xmlns="https://www.dmg.org/PMML-4_4"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">

  <Header copyright="Example.com"/>
  <DataDictionary> ... </DataDictionary>

  ... a model ...

</PMML>
```

Opensccoring

Overview Products Licensing Consulting Blog

Standards-based, Open-source Middleware for Predictive Analytics Applications

Convert your fitted **Scikit-Learn**, **R** or **Apache Spark** models and pipelines into the standardized **Predictive Model Markup Language** (PMML) representation, and make quick and dependable predictions in your **Java/JVM** application.

Standardization enables automation, which in turn enables higher efficiency and higher quality business processes.

GET STARTED NOW!

Villu Russmann and his team have developed tools to make PMML work from the most commonly used advanced Python and Spark algorithms/models

https://openscoring.io/

Data preparation is also part of a Deployment pipeline

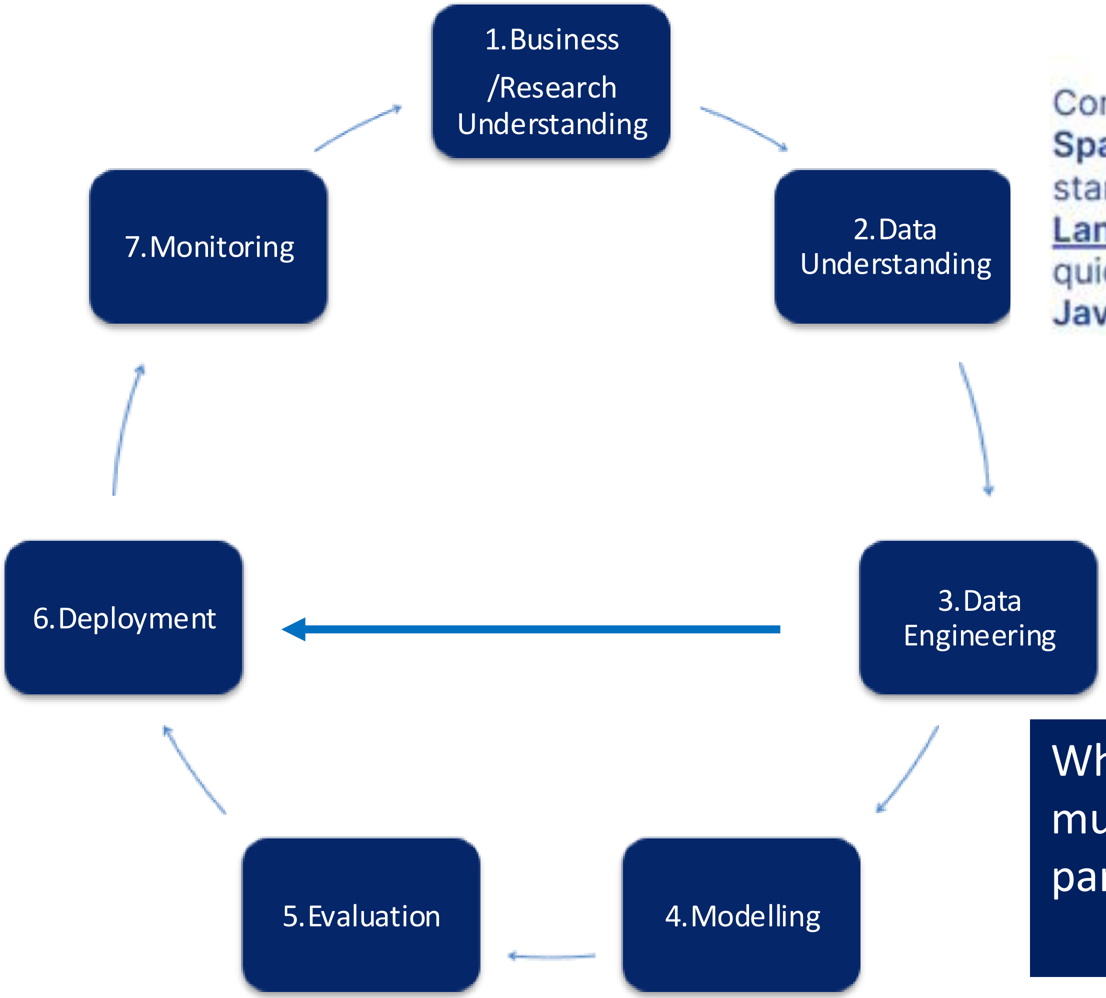

Convert your fitted **Scikit-Learn**, **R** or **Apache Spark** models and pipelines into the standardized **Predictive Model Markup Language** (PMML) representation, and make quick and dependable predictions in your **Java/JVM** application.

What happens in **Data Engineering** must also happen in **Deployment** as part of the broader pipeline

An Example Case

- In a recent project, we worked to develop a model to forecast operational customer satisfaction for a UK bank
- We tested multiple algorithms, including **ARIMAS** and various **Regressions**
- **XGBoost** was the most accurate
- The bank needed to deploy the model in a forecasting (What-if?) simulator in Excel
- In the end, we were not able to install any additional open-source engines … for technical reasons
- Hence, the currently deployed model uses a combination of **Factor Analysis and Constrained Linear Regression**
    - – The accuracy is still acceptable, but we had to compromise it
- The Openscoring team has developed a working prototype to deploy XGBoost models into MS Excel
    - Supports Python's Pickle, R's RDS or Predictive Model Markup Language (PMML)

OpenScoring's Excel XLSBoost deployer

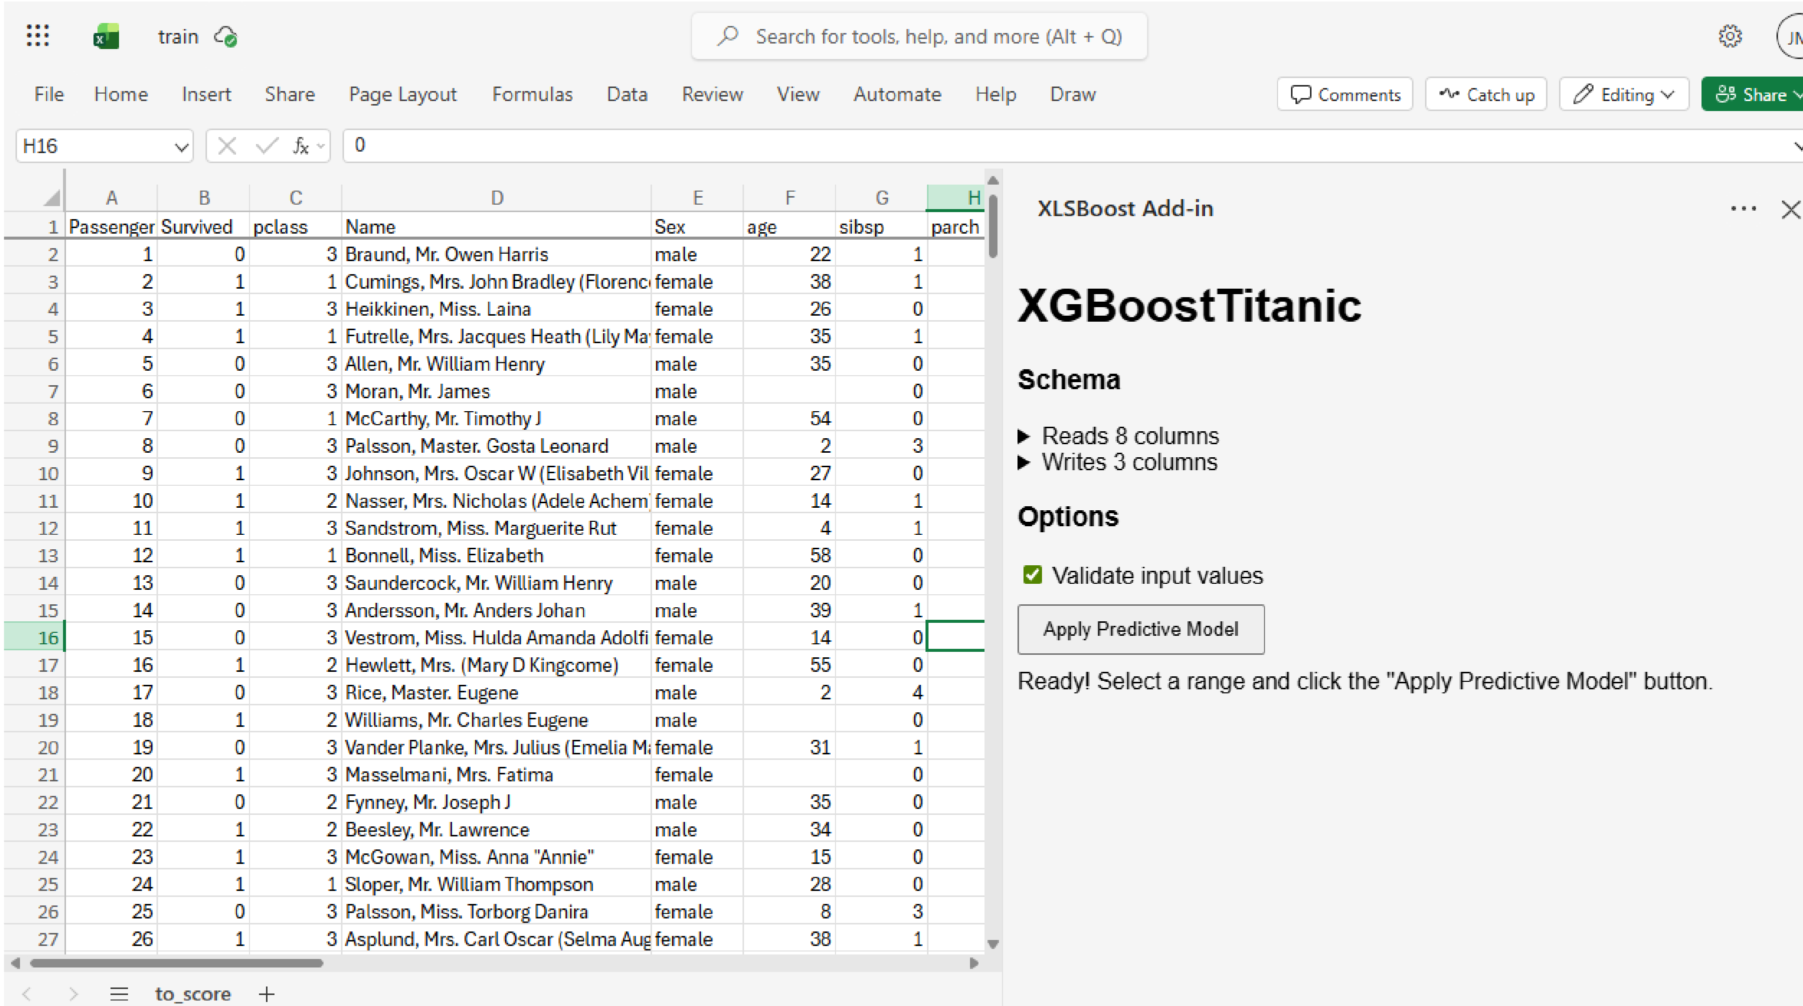

	A	B	C	D	E	F	G	H
1	Passenger	Survived	pclass	Name	Sex	age	sibsp	parch
2	1	0	3	Braund, Mr. Owen Harris	male	22	1	
3	2	1	1	Cumings, Mrs. John Bradley (Florenc	female	38	1	
4	3	1	3	Heikkinen, Miss. Laina	female	26	0	
5	4	1	1	Futrelle, Mrs. Jacques Heath (Lily Ma	female	35	1	
6	5	0	3	Allen, Mr. William Henry	male	35	0	
7	6	0	3	Moran, Mr. James	male		0	
8	7	0	1	McCarthy, Mr. Timothy J	male	54	0	
9	8	0	3	Palsson, Master. Gosta Leonard	male	2	3	
10	9	1	3	Johnson, Mrs. Oscar W (Elisabeth Vil	female	27	0	
11	10	1	2	Nasser, Mrs. Nicholas (Adele Achem	female	14	1	
12	11	1	3	Sandstrom, Miss. Marguerite Rut	female	4	1	
13	12	1	1	Bonnell, Miss. Elizabeth	female	58	0	
14	13	0	3	Saundercock, Mr. William Henry	male	20	0	
15	14	0	3	Andersson, Mr. Anders Johan	male	39	1	
16	15	0	3	Vestrom, Miss. Hulda Amanda Adolfi	female	14	0	
17	16	1	2	Hewlett, Mrs. (Mary D Kingcome)	female	55	0	
18	17	0	3	Rice, Master. Eugene	male	2	4	
19	18	1	2	Williams, Mr. Charles Eugene	male		0	
20	19	0	3	Vander Planke, Mrs. Julius (Emelia Ma	female	31	1	
21	20	1	3	Masselmani, Mrs. Fatima	female		0	
22	21	0	2	Fynney, Mr. Joseph J	male	35	0	
23	22	1	2	Beesley, Mr. Lawrence	male	34	0	
24	23	1	3	McGowan, Miss. Anna "Annie"	female	15	0	
25	24	1	1	Sloper, Mr. William Thompson	male	28	0	
26	25	0	3	Palsson, Miss. Torborg Danira	female	8	3	
27	26	1	3	Asplund, Mrs. Carl Oscar (Selma Aug	female	38	1	

https://xlsboost.com/

Option 3 – Using the same engine to build (train) and deploy

The basic idea

A. Start with the modelling “pipeline/job”
   - You typically need to adapt it to create a deployment/scoring version
     - Usually a simplified version
B. Install the software on a deployment machine; R, Python, SAS, SPSS, WPS, etc.
   - Or use the same machine
C. Periodically run new data through the job and send the scores/predictions to an operational target

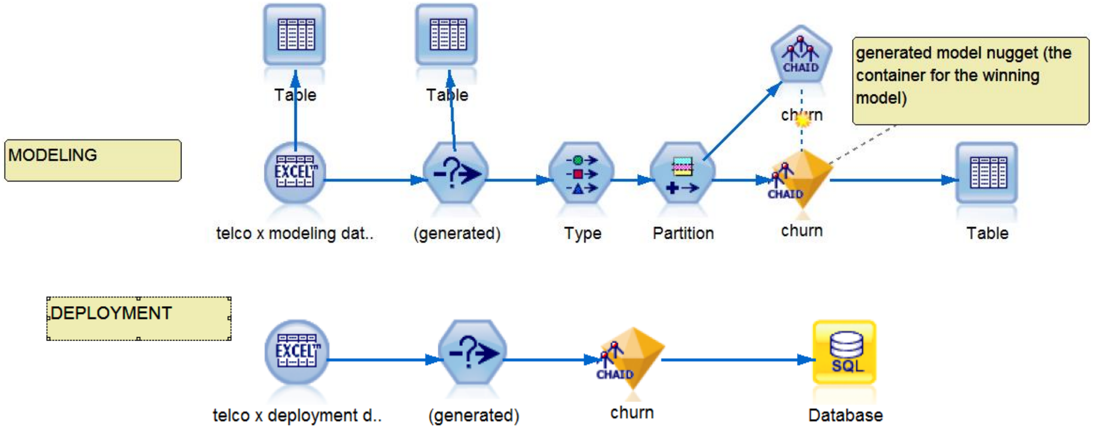

Which means we can usually <u>schedule</u> easily

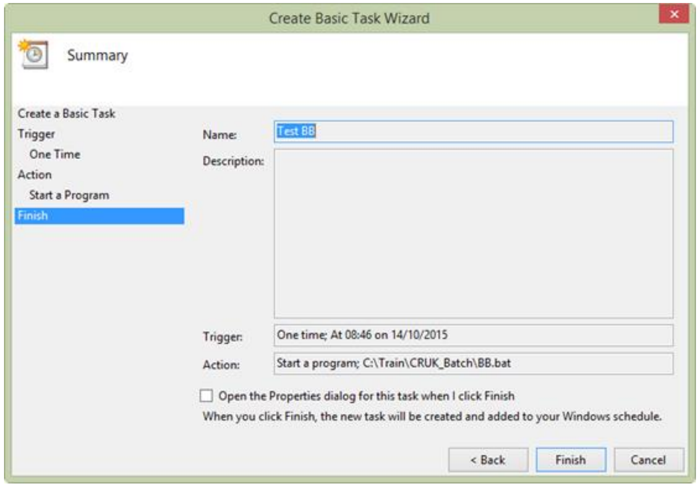

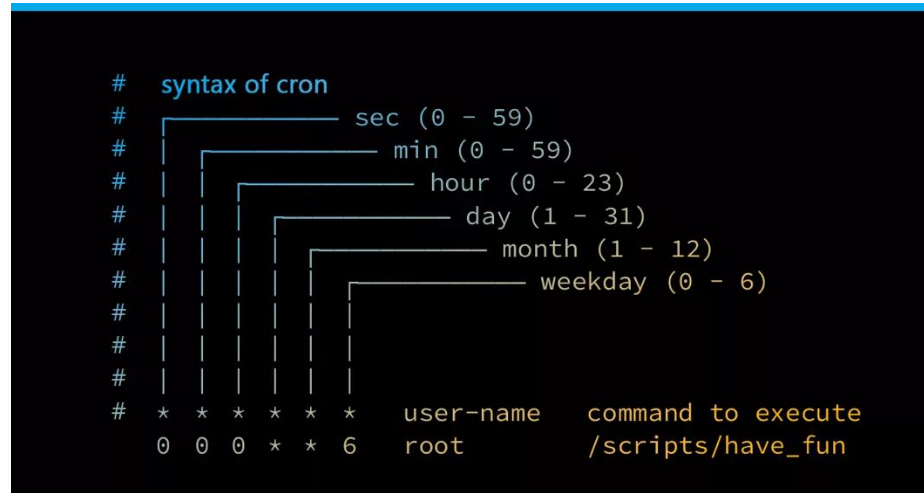

As long as the analytical tool/engine allows **batch execution** (usually via the command line) we can use standard scheduling tools in the OS e.g. Windows Task Scheduler or CRON to run at a given time/periodically

Real-time would require more engineering

Option 4 – Use a Data Science/ML platform

Platform essentials

Support for low/no-code
- Visual interfaces
- May optionally generate code

And code
- Scripts
- Noteboooks

Support collaboration
- Roles
- Within functions
- Between functions

Auto ML
- Algorithm selection
- Algorithm settings optimisation
- Feature selection

Ease of deployment
- Scheduled (batches)
- Real-time

Post deployment life cycle managment
- Monitoring
- Reporting/Alerting
- Model refreshing
- Model retirement

MLOps support
- Enables integration and scaling of deployments across operational environments e.g. Digital, CRM, Finance, Sales, Contact centre, etc.

Platform nice-to-haves

Recommendations

- Auto ML +
- Data engineering
  - Feature extraction/creation

Advanced stuff

- Simulation
- Optimisation
- Deep learning

SDKs and APIs

- Support more programmatic code-based model development
  - C, Java, etc.
- API into platform

Model interpretability

- Tools to help understand models better
  - Many algorithms product opaque models

Support for GenAI

- Access to LLMs
  - E.g. for RAG
- Agentic support
- Foundation models

Magic Quadrant for Data Science and Machine Learning Platforms

[Gartner MQ 2025]

- DS/ML Specialists
  - Dataiku, DataRobot, DataBricks, SAS, Altair, H2O
- Broader Cloud Platforms
  - Microsoft, Amazon, Google, IBM, Alibaba
- Open source alternatives exist beyond Anaconda e.g. Airflow and MLOps

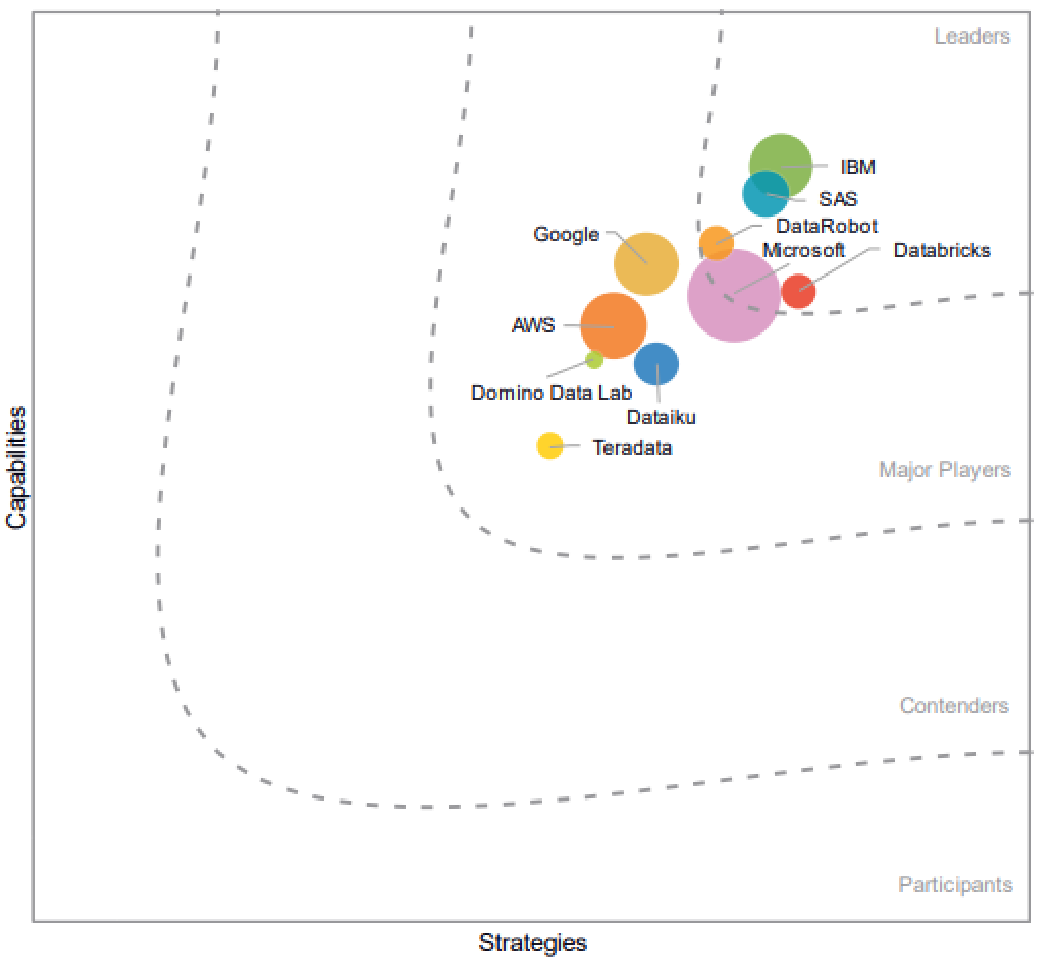

IDC MarketScape Worldwide Machine Learning Operations Platforms Vendor Assessment

[IDC Marketplace 2024]

Canvass for Low/No-code

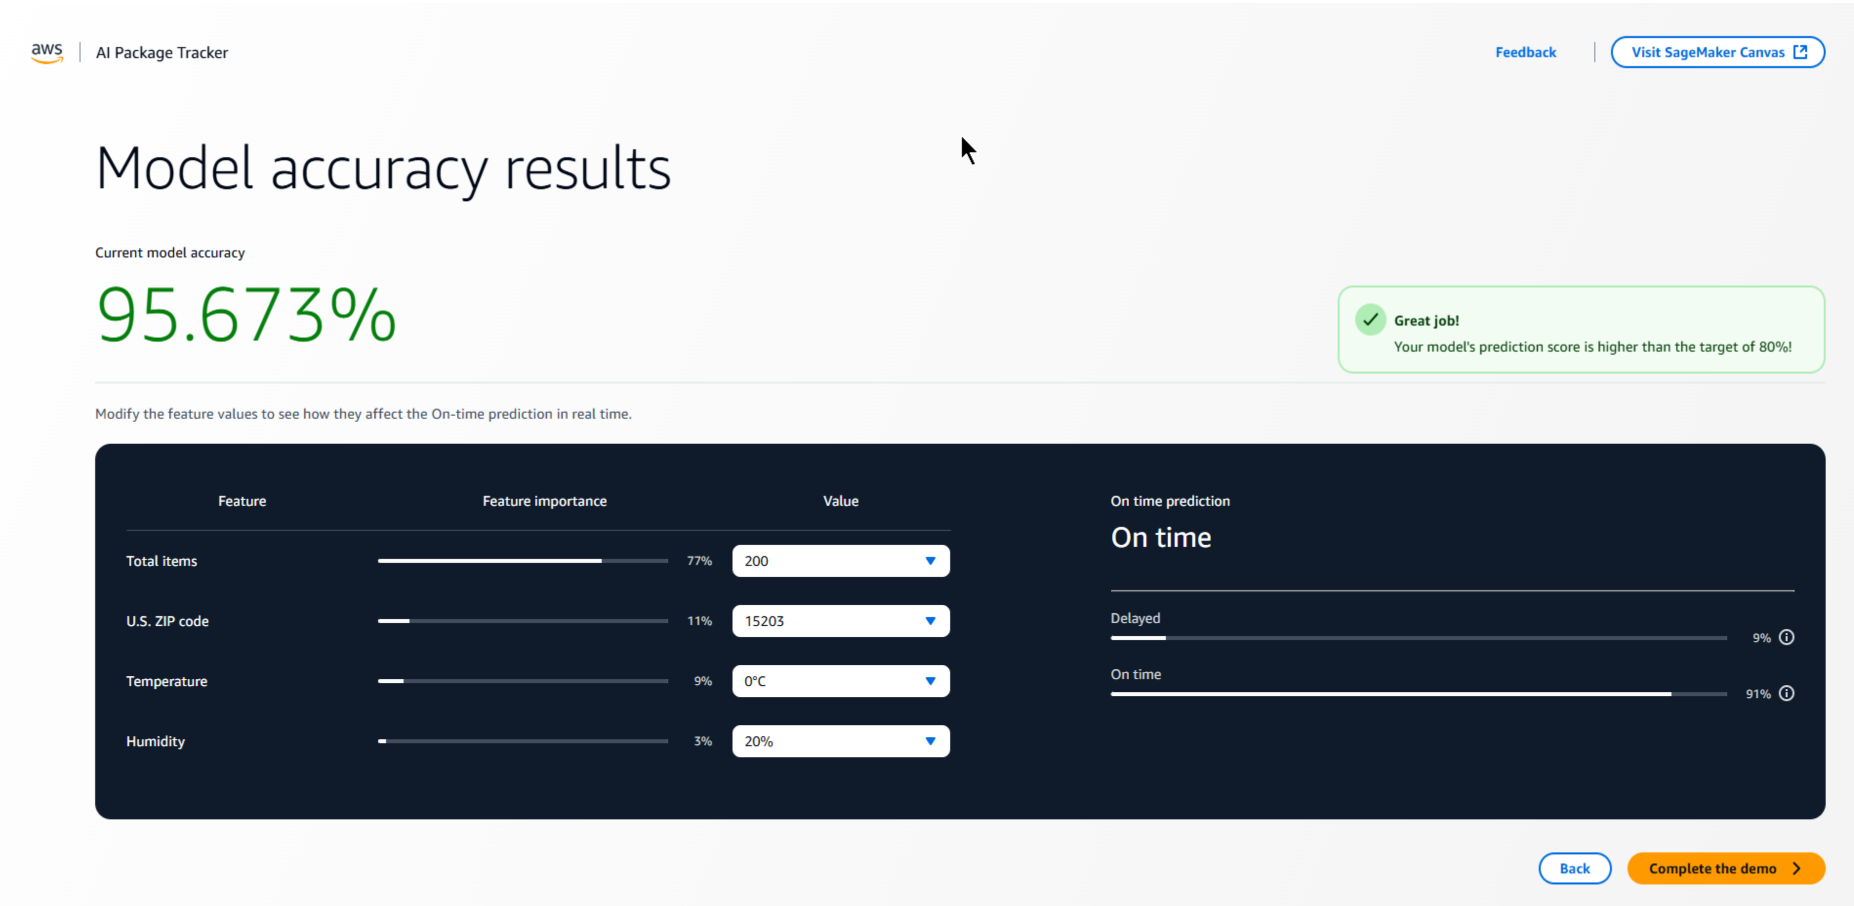

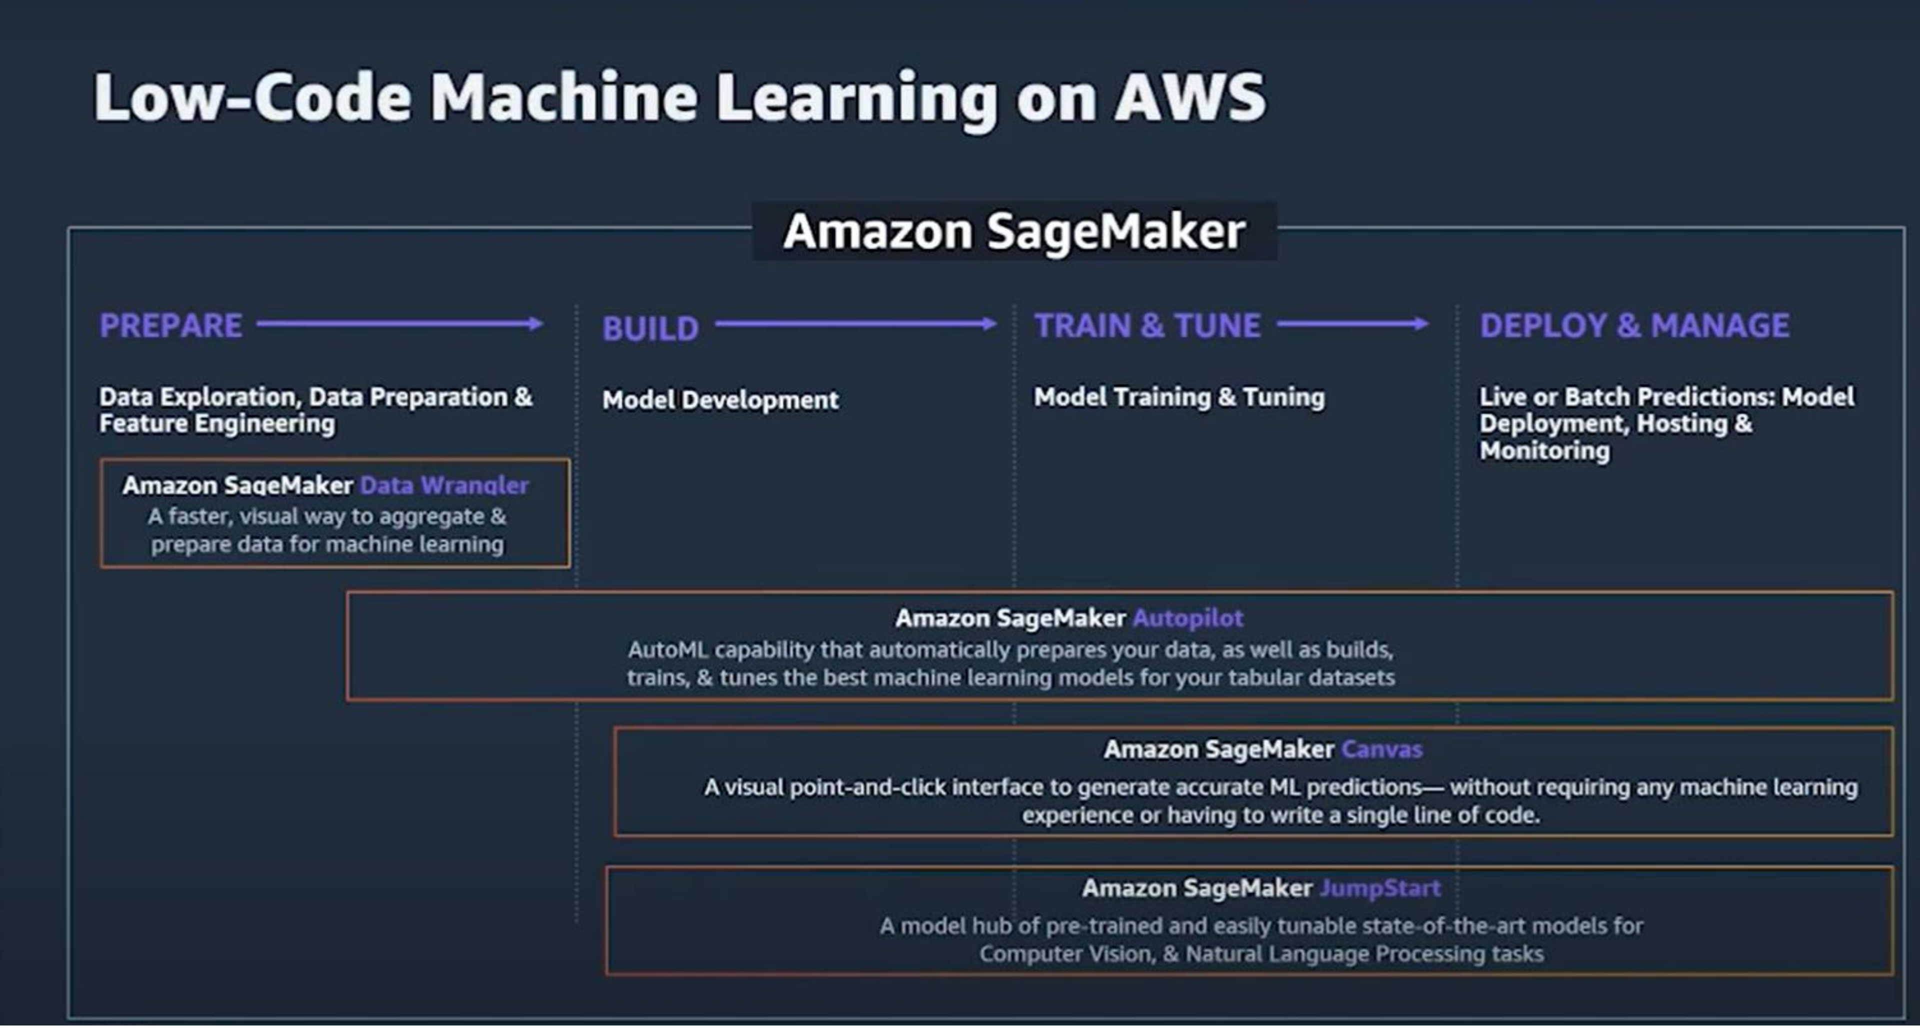

AWS Sagemaker support Low-Code ML through to deployment and monitoring

Notebooks etc. for Code

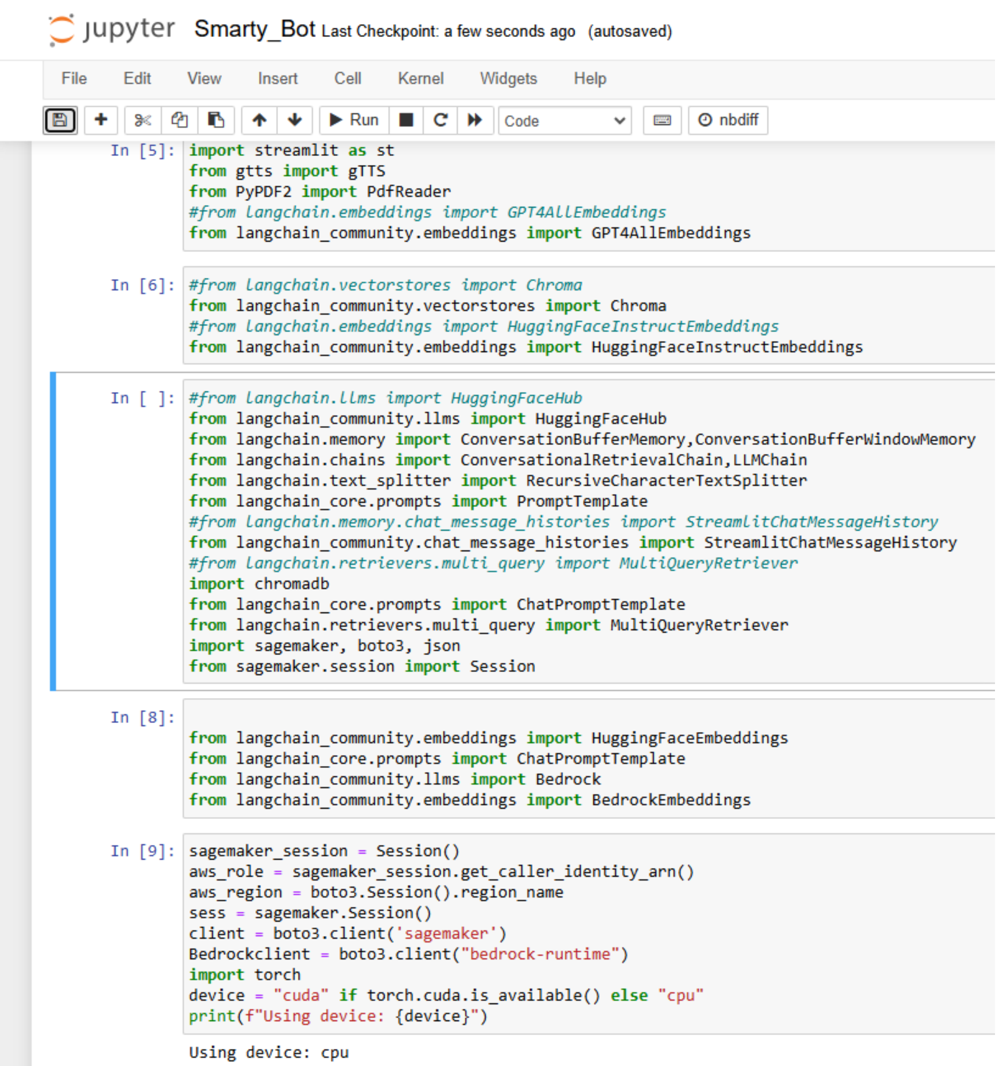

Real-time deployment from code

Hosting / Inference

Once the training is done, we can deploy the trained model as an Amazon SageMaker real-time hosted endpoint. This will allow us to make predictions (or inference) from the model. Note that we don't have to host on the same type of instance that we used to train. Because instance endpoints will be up and running for long, it's advisable to choose a cheaper instance for inference.

```
In [23]: text_classifier = bt_model.deploy(initial_instance_count = 1,instance_type = 'ml.m4.xlarge')
         ---------------------------------------------------------------------------------------!
```

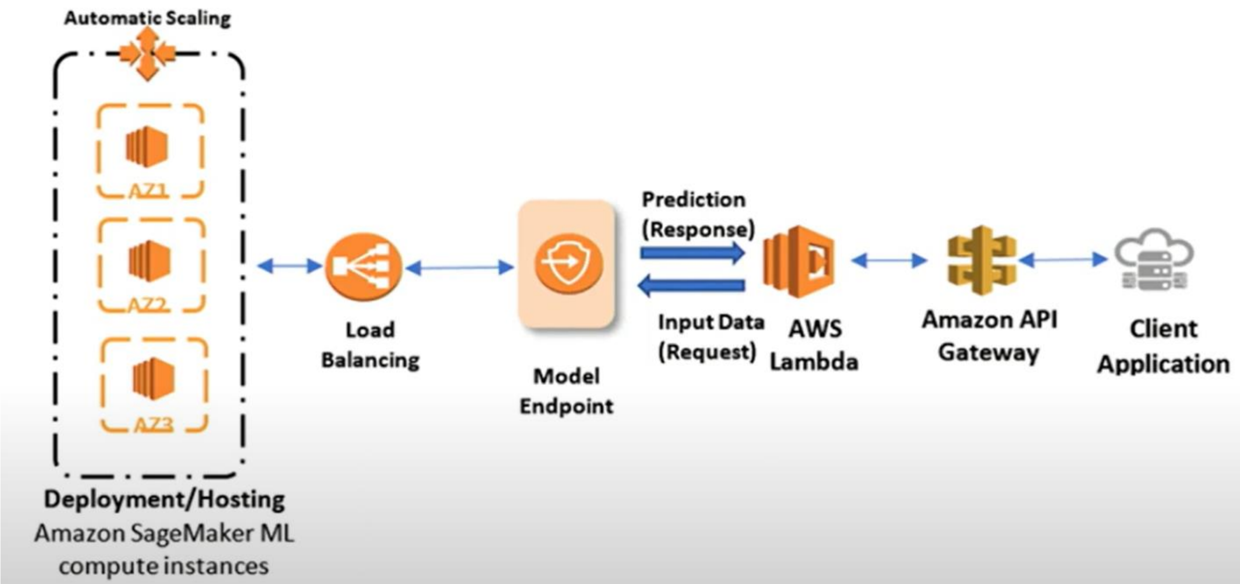

Premised platforms are also available

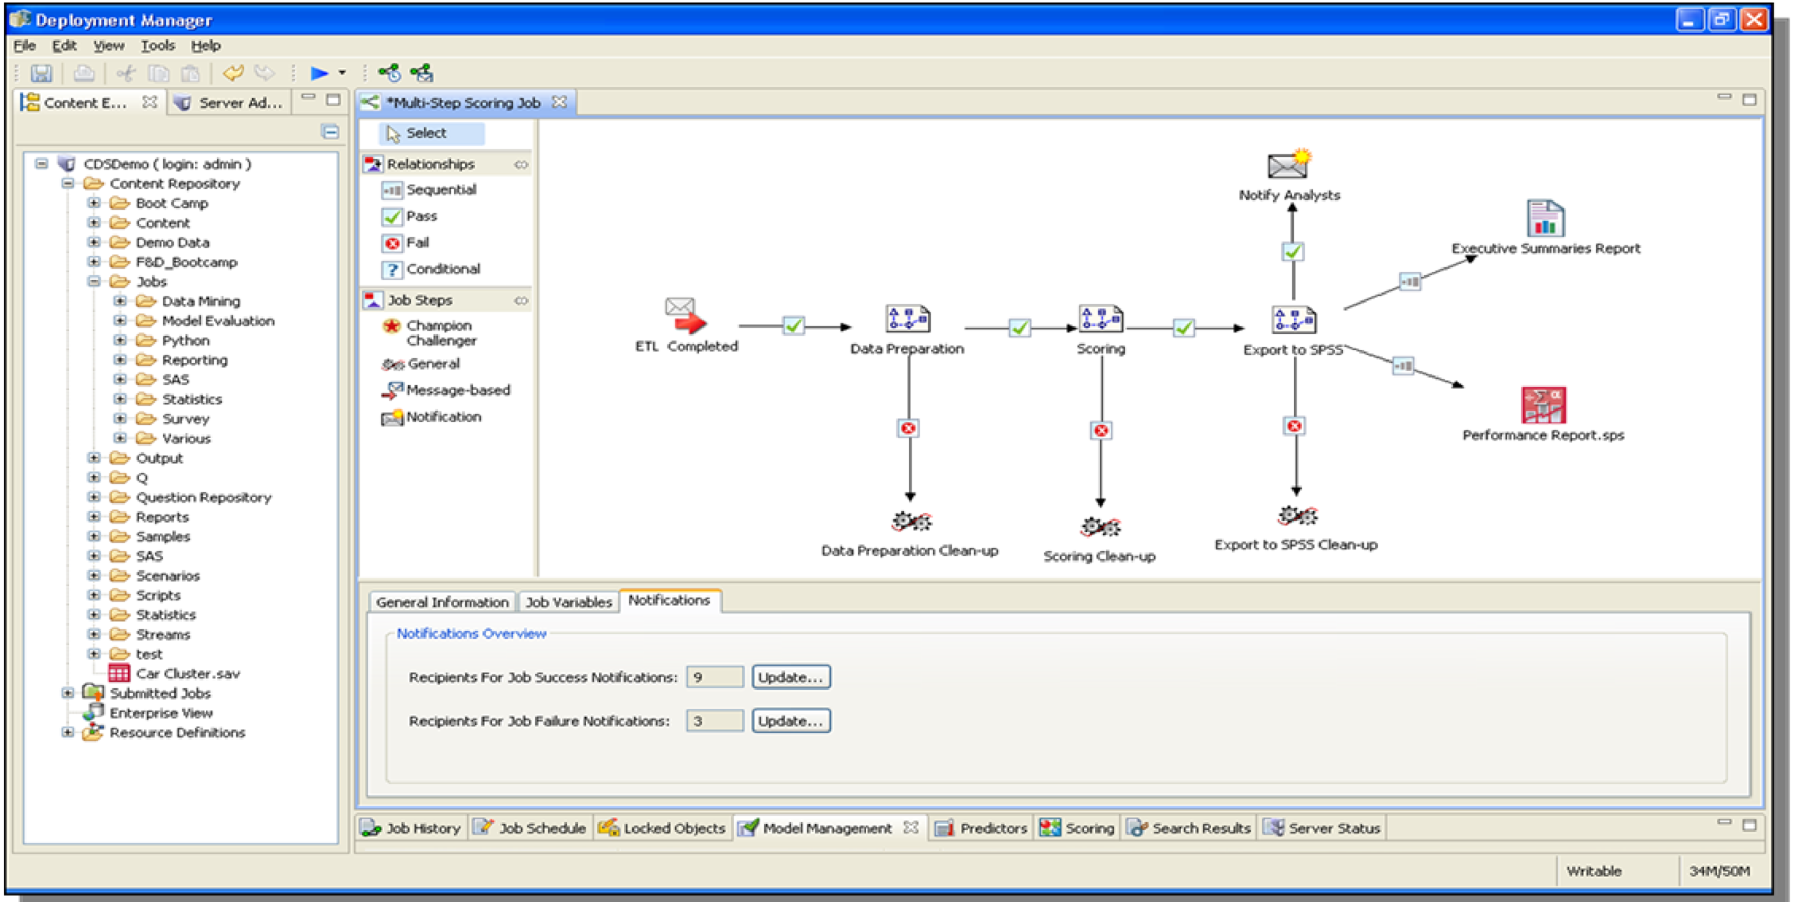

IBM/SPSS Collaboration & Deployment Services

Premised platforms are also available

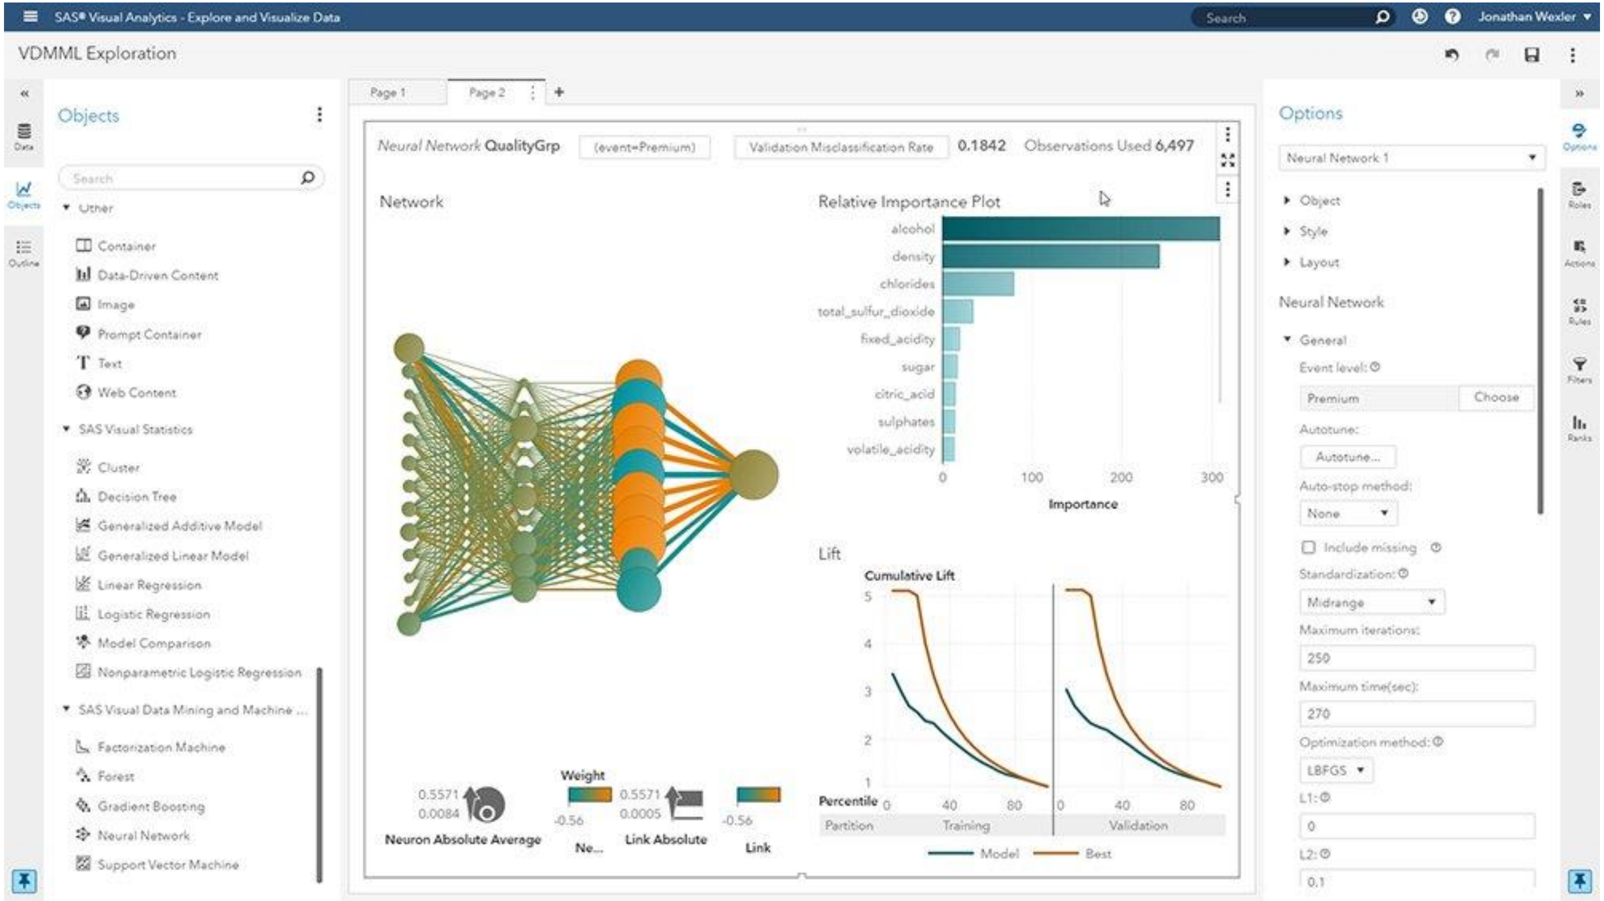

SAS Viya is cloud native but it can be installed on-prem using containers

Premised platforms are also available

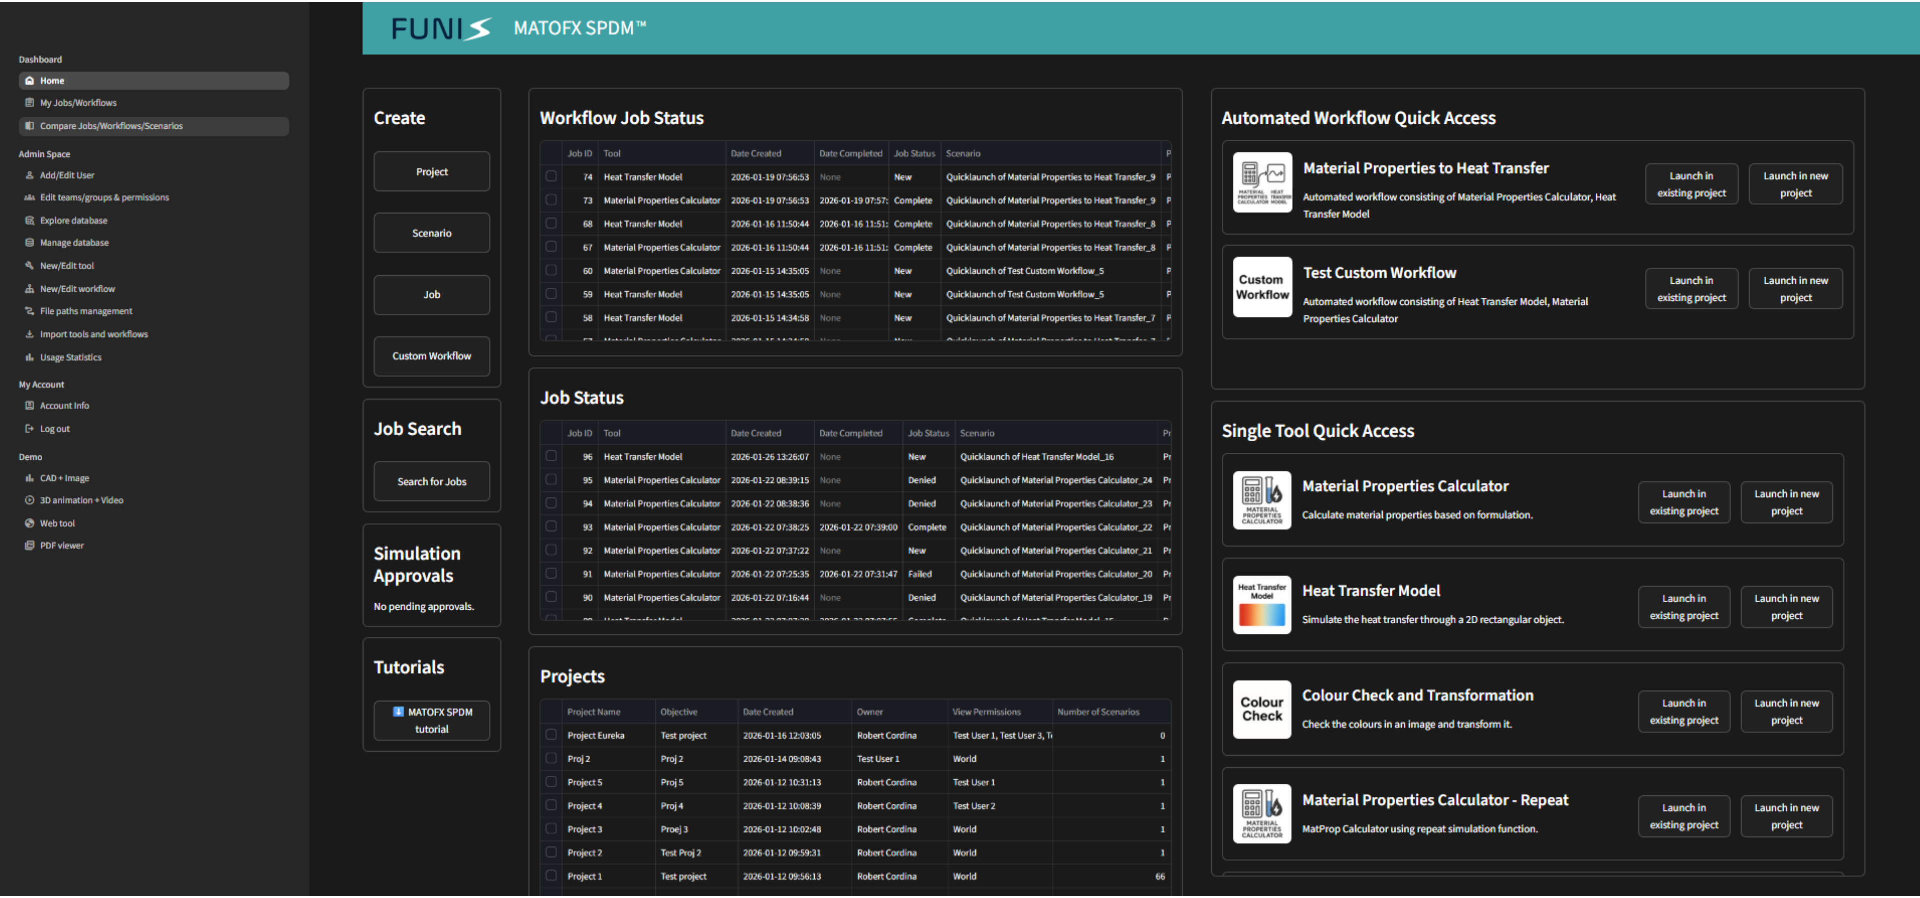

[MATOFX SPDM](#) **from Funis Consulting** is a tool agnostic simulation workflow platform

IBM Watson X - Auto AI RAG Experiment progress map

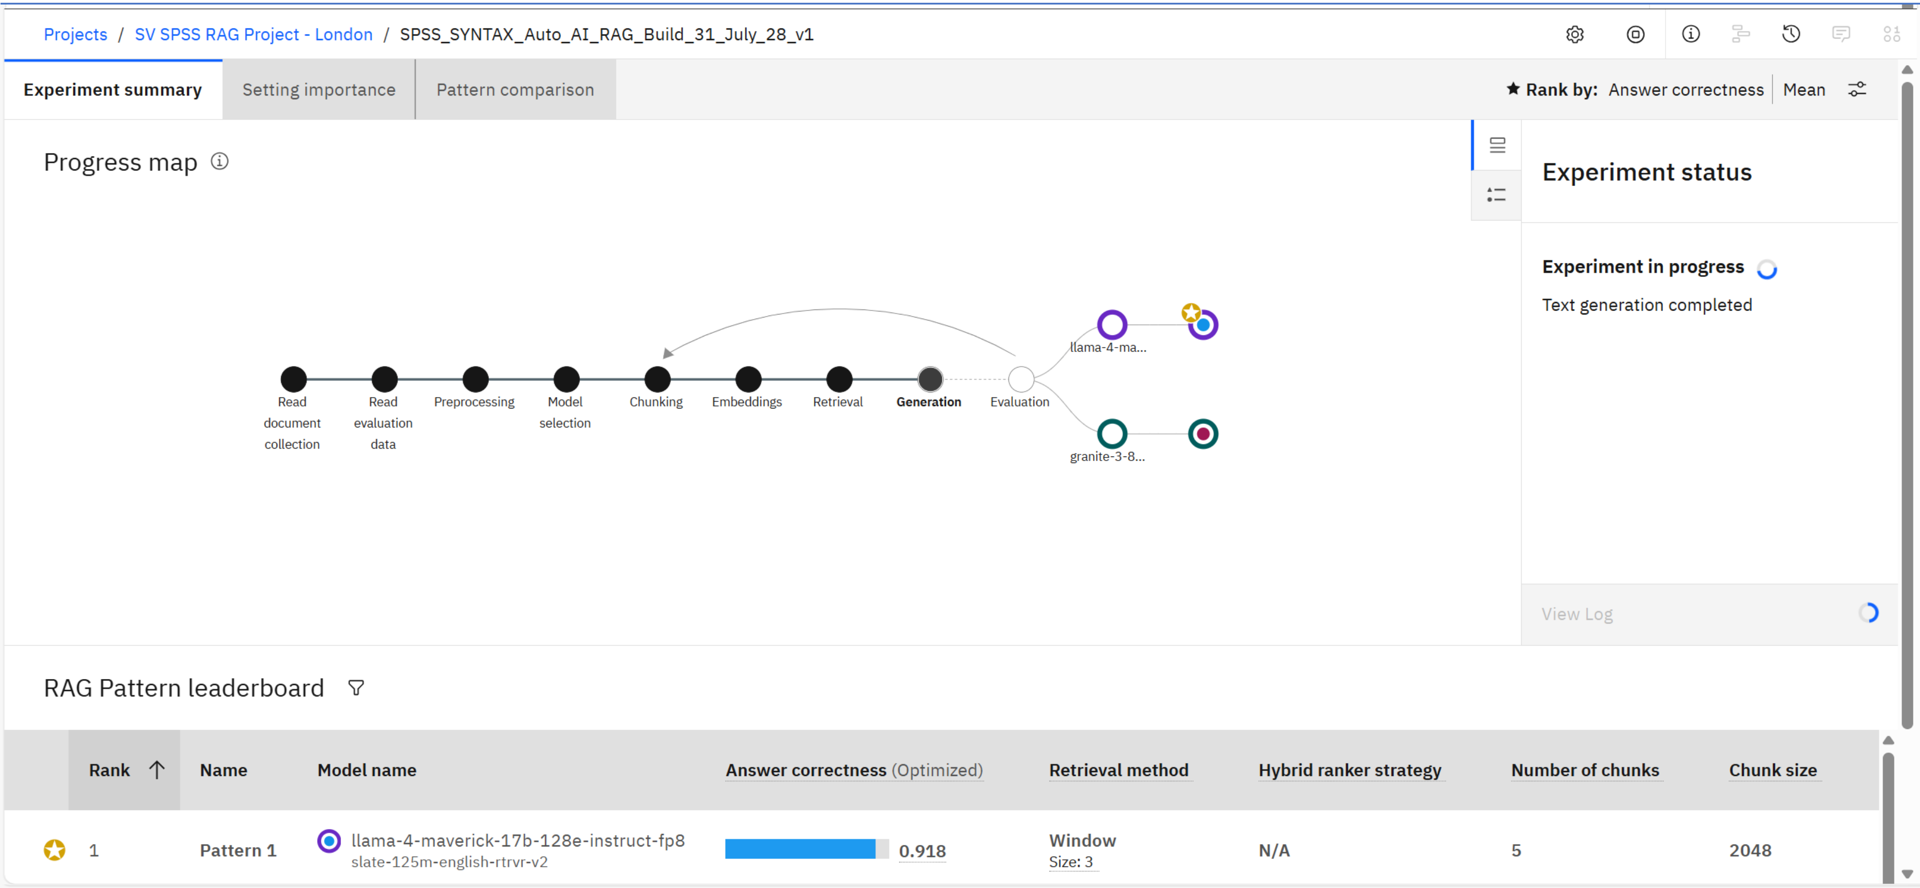

IBM Watson X - API endpoint of deployed AI Service

chatspss.com

Matilda LIVE | Google Sheets | Matilda DEV | Thane Spa DEV | (3) WhatsApp | Cloud Computing S... | Form Recognizer St... | godaddy WordPress... | All Bookmarks

CONVERSATIONS

- Can SPSS run SARIMA?
  - Can SPSS run SARIMA

MODEL
main model

+ SEND FEEDBACK

EXPORT PDF | LOG OUT | PRIVACY | SMARTVISION

Can SPSS run SARIMA?

Here are the main points to answer the question:

- IBM SPSS Statistics can estimate SARIMA models using the ARIMA command.
- The ARIMA command supports both nonseasonal and seasonal univariate ARIMA models.
- The MODEL subcommand is used to specify the SARIMA model, including the order of the autoregressive, differencing, and moving average components, as well as the seasonal components [S2].

To access the ARIMA procedure in SPSS, you can use the following menu path: *Analyse ▸ Forecasting ▸ Create Models*

Here's an example of how to specify a SARIMA model using the ARIMA command in SPSS Syntax:

```
ARIMA SALES
  /MODEL=(0,1,1)(0,1,1).
```

This example estimates a SARIMA(0,1,1)(0,1,1) model for the `SALES` series [S2].

DOCUMENT REFERENCES

- IBM_SPSS_Statistics_Core_System_User_G...
- IBM_SPSS_Forecasting.pdf
- IBM_SPSS_Statistics_Command_Syntax_Re...
- Python_Reference_Guide_for_IBM_SPSS_S...

https://chatspss.com/

new conversation

Blocker	Mitigations
Technical	Involve technical colleagues early Figure out the technical path to get there Automate that
Data	Improving data quality is a broad project often leading to a Data Lake/Warehouse to support Data Science Tactically assess quality and look to clean what we can and only use data of Assess if this will be sufficient to hit our targets
Stakeholder buy-in	To get alignment we need to agree what a good outcome is Agree the criteria in a less technical way e.g. Lift, gain, profit Keep them involved/informed where possible
End user buy-in	Engage early. Get feedback. Make sure there are benefits e.g. productivity, and communicate them
Model does not work in training. It isn't accurate enough or other success criteria are not hit	Need to clearly agree what "work" means with stake holders. Is any model better than no model? This is the "art+science" of building ML models which can be solved with more/better data, more feature engineering, algorithmic choices and tuning
Model works in training but not in deployment	Add a QA process Peer reviews

In summary

- There are many ways to deploy
- Deployment planning is key
  - Make sure we have a plan to deploy assuming for model success
  - Make sure we set expectations … we don’t always get to where we want
- Get communication/alignment with stakeholders early and often … don’t forget end user stakeholders who may be a broad group who become involved post deployment
- Involve technical colleagues early too
- Depending on where you are in your deployment journey think about the next level of automation
  - The more automated the less work … the less likely errors
- Platforms are definitely not for all use cases/contexts
  - Depends on scale and availability
  - Platforms/automation do not eliminate blockers on their own … but they help
- Did we say plan deployment from the beginning?

Smart Vision Service/Product Offerings

- **CONSULTING SERVICES**

  Providing expert guidance and strategic advice to help organizations achieve their analytical goals

- **IMPLEMENTATION SERVICES**

  Seamless integration and deployment of software, systems, and technologies to ensure successful adoption and utilization.

- **TRAINING, SUPPORT & ENABLEMENT**

  Comprehensive training programs and ongoing support to empower clients and ensure they maximize the value of our services. Enabling our customers to become self sufficient through collaborative projects

- **SOFTWARE LICENSING AND SUPPORT**

- **MANAGED SERVICES**

  Outsourced data, analytics and business process management solutions that optimize operations, reduce costs, and enhance efficiency.

- **CUSTOM DEVELOPMENT**

  Bespoke software and application development to address unique business/analytical requirements, increase automation and provide competitive advantages.

Contact us:

+44 (0)207 786 3568

info@sv-europe.com

X: @sveurope

Follow us on LinkedIn

Sign up for our Newsletter

Thank you