SAS6344-2016

Mass-Scale, Automated Machine Learning and Model Deployment Using SAS® Factory Miner and SAS® Decision Manager

Jonathan Wexler and Steve Sparano, SAS Institute Inc., Cary, NC

ABSTRACT

Business problems have become more stratified, and micro-segmentation is driving the need for mass-scale, automated machine learning solutions. Additionally, deployment environments include diverse ecosystems, requiring hundreds of models to be built and deployed quickly via web services to operational systems. The new SAS® automated modeling tool allows you to build and test hundreds of models across all of the segments in your data, testing a wide variety of machine learning techniques. The tool is completely customizable, allowing you transparent access to all modeling results. This paper shows you how to identify hundreds of champion models using SAS® Factory Miner, while generating scoring web services using SAS® Decision Manager. Immediate benefits include efficient model deployments, which allow you to spend more time generating insights that might reveal new opportunities, expose hidden risks, and fuel smarter, well-timed decisions.

INTRODUCTION

Business problems are increasingly complex, and the need for stratified analytics has surged. Businesses are looking to apply advanced analytic techniques to solve problems such as attrition, acquisition, and anomaly detection. Companies expect solutions that encompass data ingestion, modeling, and deployment (action), with minimal manual handoffs. Analytical problems often require advanced machine learning techniques such as random forest, neural networks, and decision trees. There is a marked shortage of skilled data scientists at many companies, so the need for automated analytics is great.

This paper shows how to use a tightly integrated solution using SAS® Factory Miner and SAS® Decision Manager to build and deploy hundreds, if not thousands, of models, all within a few clicks. You will learn how to build hundreds of models using Model Templates, which are portable analytic methods that can be shared within an organization. Once you have completed building your models, you can register your models to SAS® Model Manager, a component within SAS® Decision Manager to monitor the model performance over time. Once registered, you can deploy them easily from SAS® Decision Manager, as restful scoring services that be easily integrated with your business processes.

ANALYTICAL EVOLUTION

Data science and machine learning have exploded onto the analytic market scene. The cost for computing power has dropped to the point where companies can now invest in infrastructure that supports advanced machine learning computations. Businesses depend on analytic departments for daily decision making. Data scientists are in high demand. Unfortunately, the available skill set to build machine learning models is not sufficient at most companies. The number of models that data scientists can build is limited primarily by workload and time.

Organizations are also struggling with how to automate more of their operations to relieve the burden already on their IT staff for the deployment and monitoring of their growing analytical model inventory. As this inventory of analytical models grows, companies must invest in strategies that allow them to quickly focus their valuable and scarce resources on the value-added tasks and not on the mundane.

There is an abundance of *business* analytic talent at companies, with domain-specific knowledge of business problems. Similarly, citizen data scientists lack the ability to access and manipulate raw data and author algorithms. Building models often requires advanced feature engineering, feature selection, and model building/fine tuning.

Many software tools offer the ability for data scientists to build models, but there are not many that allow them to share *and* deploy them.

Companies that enable their data scientists to focus on building solid analytic foundations for their colleagues will be able to scale much faster. Scaling and collaboration enable the business to make decisions rapidly.

The decisions within companies are at various points throughout their business network of systems, partners, suppliers, and customers. Throughout this network are critical decision point into which the analytical models need to be injected to enable more precise and repeatable decisions, but if the processes that companies rely on are manual and error-prone, the value of putting models to work quickly is diminished.

ANALYTIC LIFECYCLE

Depending on your level of integration with the business, the analytical lifecycle might follow different paths. Primarily, we address the modeling lifecycle in two major stages. The first stage is Discovery, where you are interested in uncovering new trends, patterns, and signals in your data. This is an iterative stage in the lifecycle as many experiments might be run, looking for the best model(s) to solve your business problem. Figure 1 depicts this two-staged approach.

Deployment, the second stage of the analytical lifecycle, involves moving the analytical models and decisions into business processes. Once the analytical models are developed, the business requires an approach that seamlessly allows the models to be quickly published into business processes, monitored for optimal performance, alerted when changes in performance are detected, and to quickly move the models and decisions into a business process. The process needs the flexibility to accommodate the various targets where the decisions are made: real-time and batch, including in-database.

Figure 1. Analytics Approach Depends on the Goal – Discovery Versus Deployment

Discovery might involve multiple data sources, both structured and unstructured, to maximize the chance for a robust solution. At this stage in the process, it's ok to *fail fast.* Failing at this point in the process is a good thing. I would rather know what methods do not work early in the process, rather than finding out once the models have been deployed into production.

Ideally, at this stage, you can experiment with multiple machine learning techniques very quickly and automatically. Data scientists often chain together feature creation, feature extraction, dimension reduction, and machine learning/predictive models. Features created or extracted in the data often provide significance in models, rather than just relying on the inputs.

Flexibility with the models allows various decision scenarios to be analyzed, whether those scenarios include real-time or batch target environments. This flexibility allows the data scientists and business teams to understand the impact of new models within the decision and not separate from them.

SAS SOLUTIONS TO SUPPORT DISCOVERY AND DEPLOYMENT

SAS simplifies the discovery and deployment aspects of your work. The following section outlines a modernized process using two integrated SAS software solutions, SAS® Factory Miner and SAS® Decision Manager.

SAS® Factory Miner

SAS® Factory Miner is a new browser-based tool that runs on the third maintenance release for SAS 9.4 platform. It enables you to automatically build champion models for each of the segments in your data. Take advantage of model templates, which are interactive modeling recipes. Experiment with multiple machine learning techniques, while investigating results and modeling exceptions. The system is not black box. You can fine-tune your models, optimizing for the best lift, misclassification, Kolmogorov-Smirnov Statistics, lowest misclassification, and other statistical measures of fit. You can generate SAS score code for all models, and register models to SAS® Model Manager, a component of SAS® Decision Manager.

The following modeling algorithms are supported in SAS® Factory Miner:

- Bayesian Network
- Decision Tree
- Generalized Linear Model
- Gradient Boosting
- Neural Network
- Random Forest
- Regression
- Support Vector Machine

Users have the option to create their own model templates using the aforementioned algorithms as a baseline. These user-generated model templates are visible to other users of the system for further collaboration.

Within model templates, the following pre-processing options are available:

- Filtering
- Transformations
- Binning
- Imputation
- Variable Selection (Decision Tree, Random Forest, Supervised, Unsupervised, Sequential)
- Principal Components

SAS® Model Manager

SAS® Model Manager, a web-based analytical model lifecycle management tool, supports the ability to inventory models from various sources from SAS and other sources, monitor analytical model performance, improve models as new data is created, and deploy these new models to a variety of targets. SAS® Model Manager is an integral component to SAS® Decision Manager and facilitates combining models and rules together in decision flows. By providing the ability to inventory models of various types in a centralized location, companies can create a centralized and complete model inventory to better understand their analytical model capabilities and needs.

Model sources include the following:

- SAS® Factory Miner

- SAS® Enterprise Miner™
- PMML
- R language model files
- SAS/STAT®
- SAS® Econometric and Time Series

It is important to note that analytical models can also be imported from generic sources, thereby offering the flexibility to include various model types in a single, centralized location.

Reports are critical to understanding the changing performance of a model as data changes over time, and as the number of models in the inventory increases, model monitoring automation is critical to tracking the performance of numerous models. Model performance reports provided with SAS® Model Manager allow for monitoring various aspects of a model's performance and the flexibility to create custom reports as new performance monitoring needs arise.

The reports provide the necessary insights into how the model is performing, and when the performance degrades beyond an acceptable level (for example, Lift or K-S), the data scientists can be alerted to the need for retraining. This automation allows SAS® Model Manager to monitor how well models are performing using user-defined metrics and react with alerts and automated retraining.

SAS® Model Manager supports publishing models to various targets using SAS® Scoring Accelerator to deploy models for in-database and in-Hadoop execution and does this without requiring changes to the model score code. This eliminates the need to re-validate the model score code for each environment, reducing the risk associated with rewriting score code for each environment.

SAS® Decision Builder

In addition to deploying models directly from SAS® Model Manager for scoring execution, SAS® Decision Manager is a web-based product that enables organizations to manage data, business rules, analytical models, and optimization techniques. Integrated into a consistent interface for easier accessibility, SAS® Decision Builder enables companies to combine analytical models, business rules, and conditional logic into a decision flow that are then tested and refined for various scenarios.

Once tested, the SAS® Decision Builder supports publishing the decisions (models and/or business rules) for use in both batch applications and online transactions, thereby enabling operationalizing analytics and business rules for business processes. Automating decisions with SAS® Decision Builder provides a streamlined mechanism for controlling and monitoring the rules and processes used by your organization.

CASE STUDY: HOTEL LOYALTY PROGRAM ATTRITION

We use a case study to illustrate the powerful, yet simple process to build models by segment, and register all models to SAS® Model Manager for deployment, as well as through the SAS® Decision Builder from SAS® Decision Manager.

BUSINESS PROBLEM

A hotel chain wants to minimize loyalty program attrition. The modeling data set contains variables representing account profile information and hotel geographic attributes. The data has been masked and adjusted to eliminate any identifiable information. Their customer base is segmented into two premium categories: Platinum and Gold. Based on prior business knowledge, customers in different states have different travel patterns, so they are modeled separately. The goal of this analysis is to build and identify the best attrition models for every loyalty category and state combination.

SAS® FACTORY MINER

We use SAS® Factory Miner to build models by segment. For this example, the data exists as a SAS data set registered to the SAS® Metadata Server. Input data sets must be registered to the SAS® Metadata Server in order to be used in SAS® Factory Miner. SAS® Factory Miner supports local data sets, databases, and Hadoop data sources.

Factory Metadata

As new projects are created, Factory Metadata allows you to pre-assign role, level, order, transformations, and imputations automatically. This is a fantastic way for new users to log on to the system and be made aware of the predefined target and segment(s). To assign a variable role in Factory Metadata, select the data set from the Data Sources area, and select Define factory metadata. In Figure 2, we selected State as a segment, and Response as the Target. We also selected several variables to automatically reject. Any new project honors these assignments.

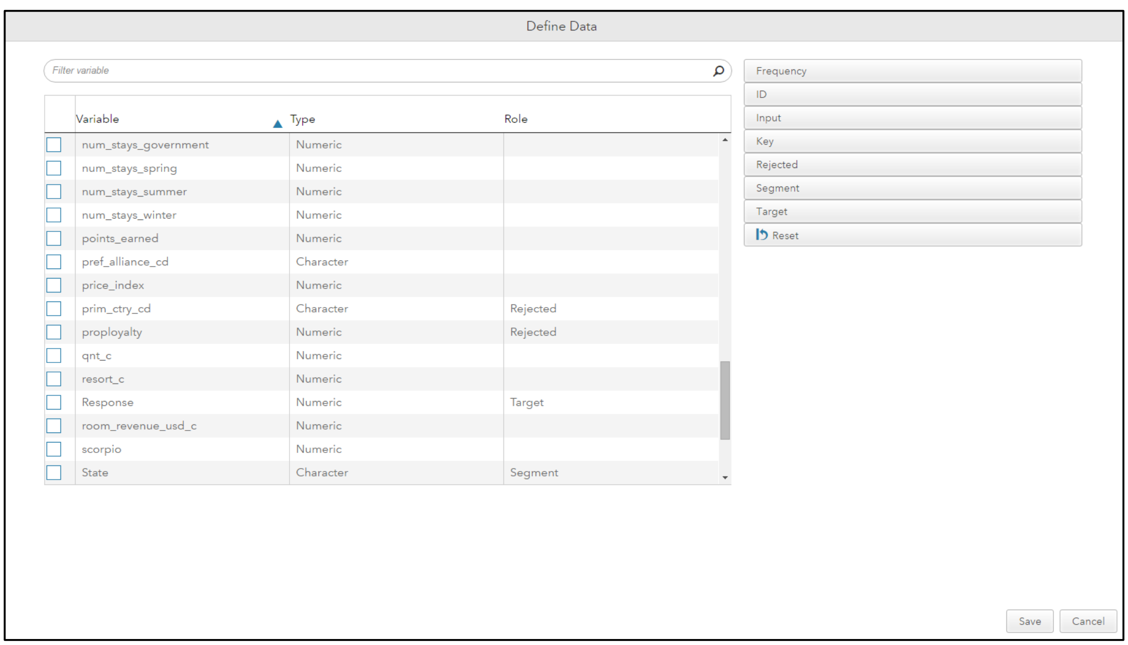

Figure 2. Factory Metadata Role Assignment: New Project

We made the process as automated as possible when creating new projects. Click the **New Project** icon from the Projects Main page. Figure 3 shows you how to name your project and select a data source. Only data sources registered to SAS® Metadata Server are shown. To ease reusability, the target and segment variables are automatically identified as you select different data sources. The system uses Factory Metadata to automatically assign roles and levels. The target is binary (Attrition=1 / Attrition=0).

By default, we partition the data into 70% training and 30% validation. You have the option to turn this off by default in account settings. It is recommended to use partitioning so that your models generalize to holdout populations. SAS® Factory Miner models automatically assess your models against the validation set. Once you click **Save**, the project is officially created.

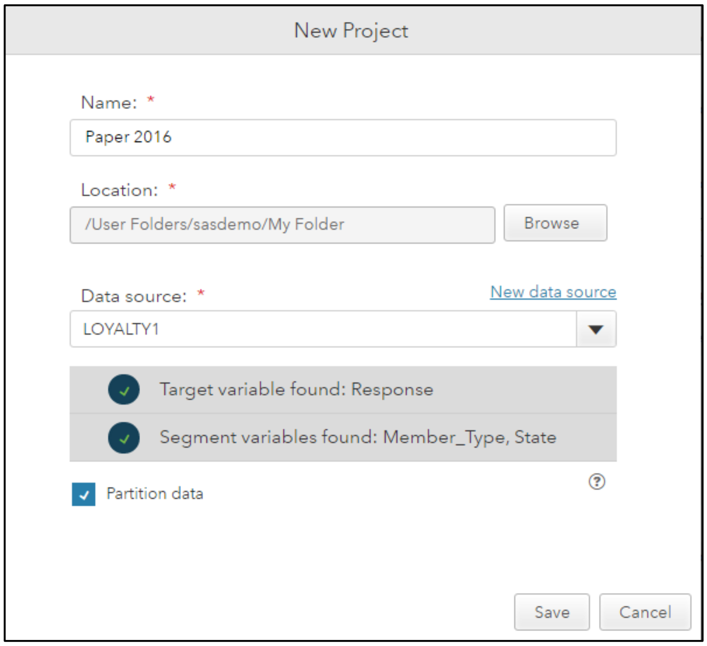

Figure 3. Create a New Project

Projects can have only one target and up to three segment variables. Warnings are displayed if you select more than one target or more than three segment variables.

Data

Once you click **Save**, SAS® Factory Miner automatically scans your input data set to intelligently assign roles, levels, transformations, and imputations. For example, variables with a limited number of distinct values are assigned as nominal. Variables that have all identical values are assigned unary, and then rejected. You have the ability to change roles and levels based on your model criteria.

In Figure 4, the variables pref_alliance_cd and prim_ctry_cd were automatically rejected since they were assigned as rejected by Factory Metadata.

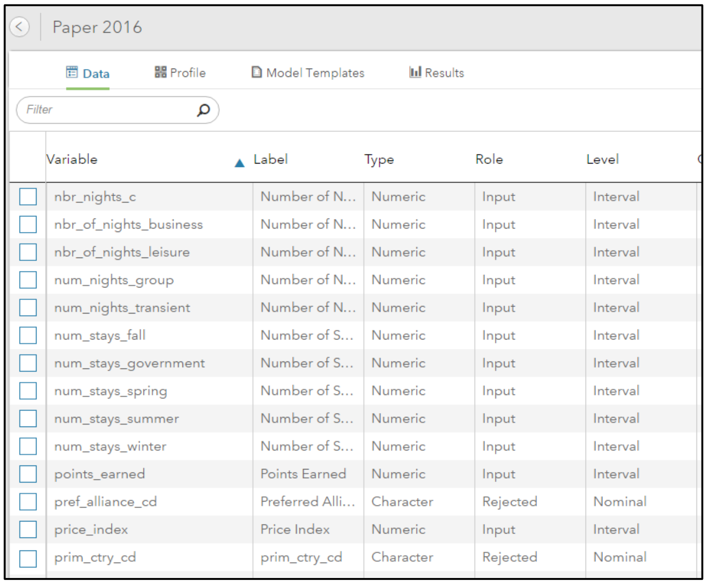

Variable	Label	Type	Role	Level
nbr_nights_c	Number of N...	Numeric	Input	Interval
nbr_of_nights_business	Number of N...	Numeric	Input	Interval
nbr_of_nights_leisure	Number of N...	Numeric	Input	Interval
num_nights_group	Number of N...	Numeric	Input	Interval
num_nights_transient	Number of N...	Numeric	Input	Interval
num_stays_fall	Number of S...	Numeric	Input	Interval
num_stays_government	Number of S...	Numeric	Input	Interval
num_stays_spring	Number of S...	Numeric	Input	Interval
num_stays_summer	Number of S...	Numeric	Input	Interval
num_stays_winter	Number of S...	Numeric	Input	Interval
points_earned	Points Earned	Numeric	Input	Interval
pref_alliance_cd	Preferred Alli...	Character	Rejected	Nominal
price_index	Price Index	Numeric	Input	Interval
prim_ctry_cd	prim_ctry_cd	Character	Rejected	Nominal

Figure 4. Data Tab Variable Role Assignment

You can change target and segments *until* the project has been initially run. Once a single model has been run, the target and segments are not customizable.

You should then click **Build Profile** to generate the segments used in the analysis.

Profile

The Profile section allows you to select which segments will be modeled in your project. The table contains the training event rate for every segment. In Figure 5, segments are automatically excluded based on project default settings. You can use the faceted filters to reduce the number of segments visible, and then click the **Exclude** or **Include** button to change your modeling set.

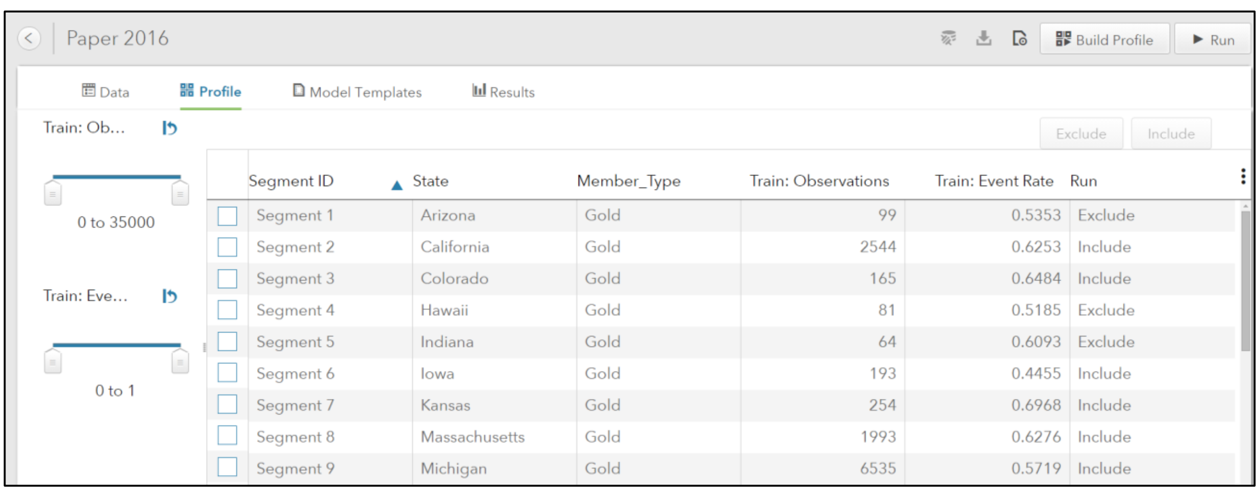

Segment ID	State	Member_Type	Train: Observations	Train: Event Rate	Run
Segment 1	Arizona	Gold	99	0.5353	Exclude
Segment 2	California	Gold	2544	0.6253	Include
Segment 3	Colorado	Gold	165	0.6484	Include
Segment 4	Hawaii	Gold	81	0.5185	Exclude
Segment 5	Indiana	Gold	64	0.6093	Exclude
Segment 6	Iowa	Gold	193	0.4455	Include
Segment 7	Kansas	Gold	254	0.6968	Include
Segment 8	Massachusetts	Gold	1993	0.6276	Include
Segment 9	Michigan	Gold	6535	0.5719	Include

Figure 5. Segment Profile

Note: You can change the segment profile *until* the project has been initially run. Once a single model has been run, the Profile is not customizable.

Model Templates

Default model templates automatically appear in this table. Default models are selected in the Model Templates section outside of a project. In Figure 6, three templates are set by default. Regression and Decision Tree are set by default. Jonathan Decision Tree is a user-generated template.

The Baseline template is automatically run in every segment so that a champion model is always chosen. The Baseline model assigns the average population target value. In this case, the segment attrition probability is assigned. We introduced the Baseline template so that when scoring models, every record is assigned a score. If no models converge, you want to have a score assigned for downstream production systems.

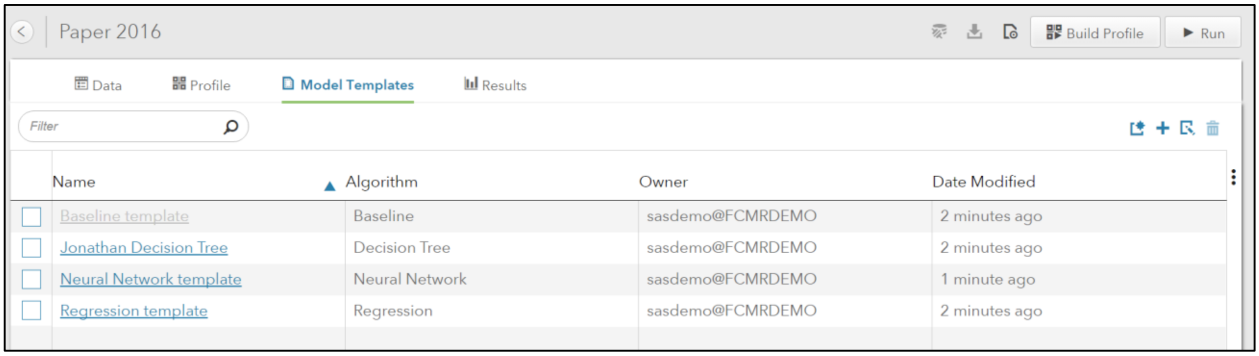

Figure 6. Model Templates within a Project

You have the ability to view the properties of any model template. In Figure 7, if you select **Regression**, you have the option to use or ignore properties, simply by dragging and dropping. Within each section, you can edit the settings of each method. If you click **Save**, the change within your project will take effect. Your model template properties outside of a project are not changed.

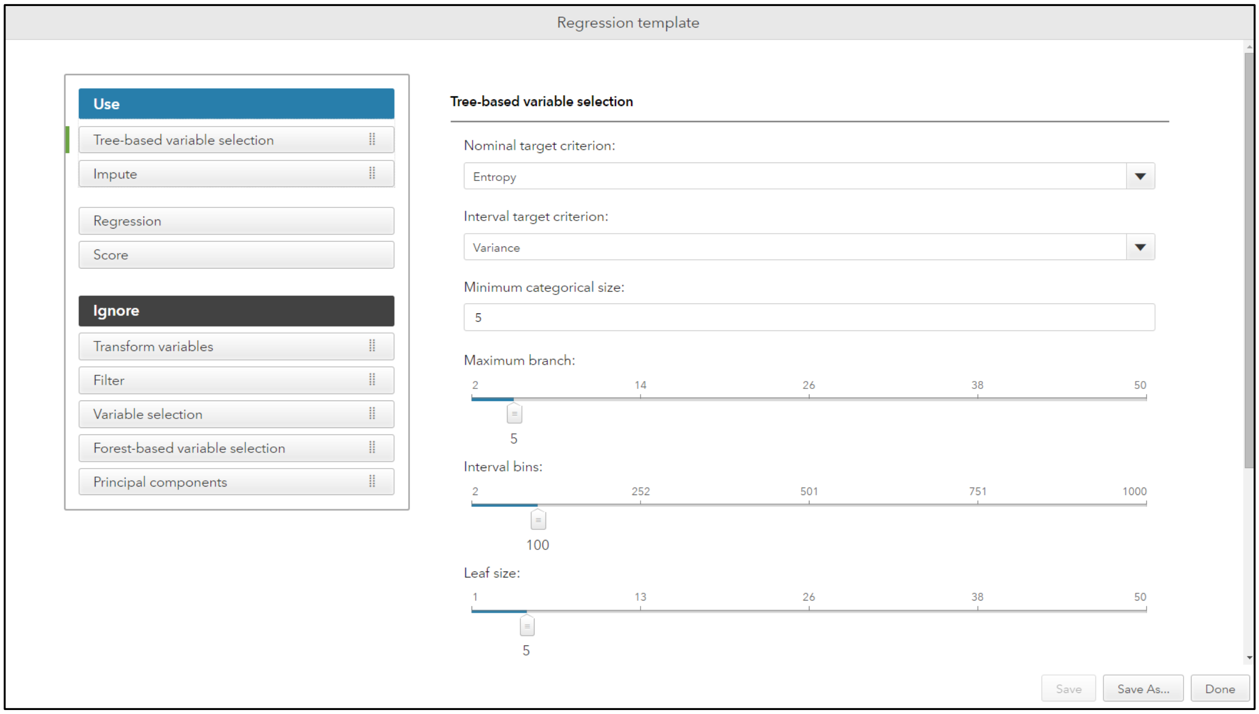

Figure 7. Edit Model Templates

From the model templates screen, click the **+** button to add in user-generated model templates.

You can delete model templates *until* the project has been initially run. Once a single model has been run, the templates cannot be deleted.

Run Project

Once you are satisfied with your data, segment, and model settings click Run. SAS® Factory Miner automatically schedules and runs models across all segments. The processing occurs in tournament fashion. Based on your hardware setup, the processing automatically occurs in parallel. Each model is an independent process, using independent processing threads. Models are run in parallel, and the results screen is updated as champion models are identified.

Champion models are selected using the default champion statistic, Validation Kolmogorov-Smirnov. As shown in Figure 8 you receive overall results across all segments, including the most popular variables chosen in champion models.

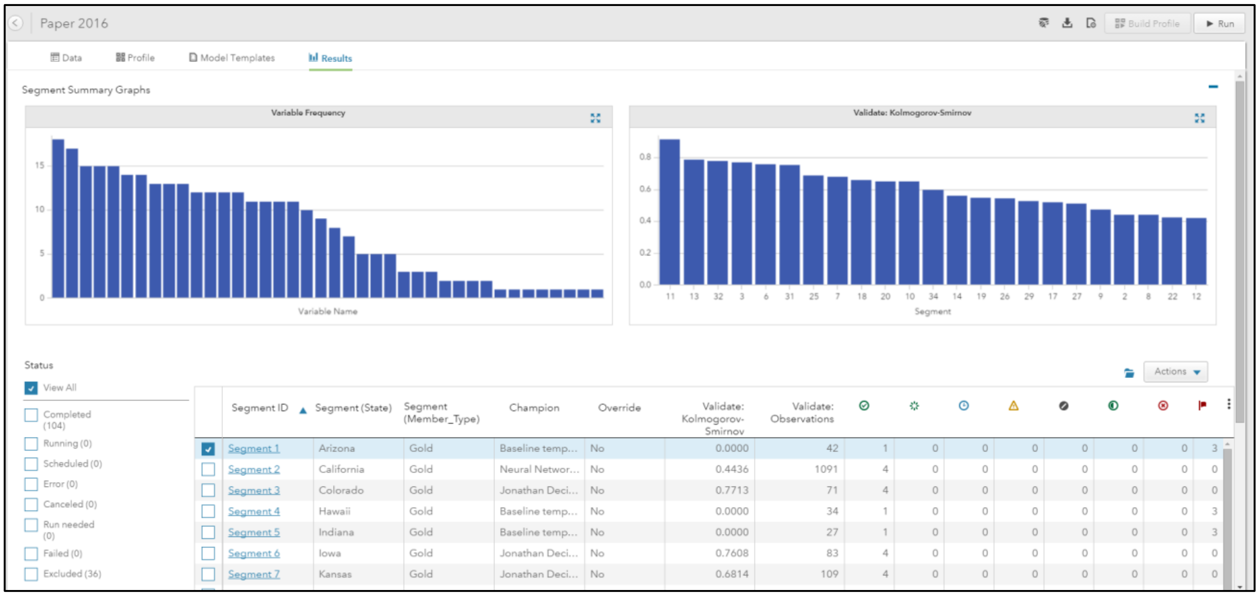

Figure 8. Project Results

You can reduce the number of segments shown on the results screen by using the faceted filter. Pay close attention to the status columns to address any segments that contain errors. Once all of your models have completed, no models will be in scheduled or running status. If you click **Run** after the segments have completed, the system will automatically run any models that are still in scheduled status.

You can safely exit the project at any point while the segments are running. All processing occurs on the SAS server via services. No processing occurs in the browser.

Results

The results are not black box! We will focus in on one segment to illustrate how to edit and improve your models. Click the **New York-Gold** segment. In Figure 9, neural network was chosen as the Champion model within the New York–Gold segment, since it has the greatest Kolmogorov-Smirnov statistic value in the validation set. As you select different models, you receive associated model results and plots.

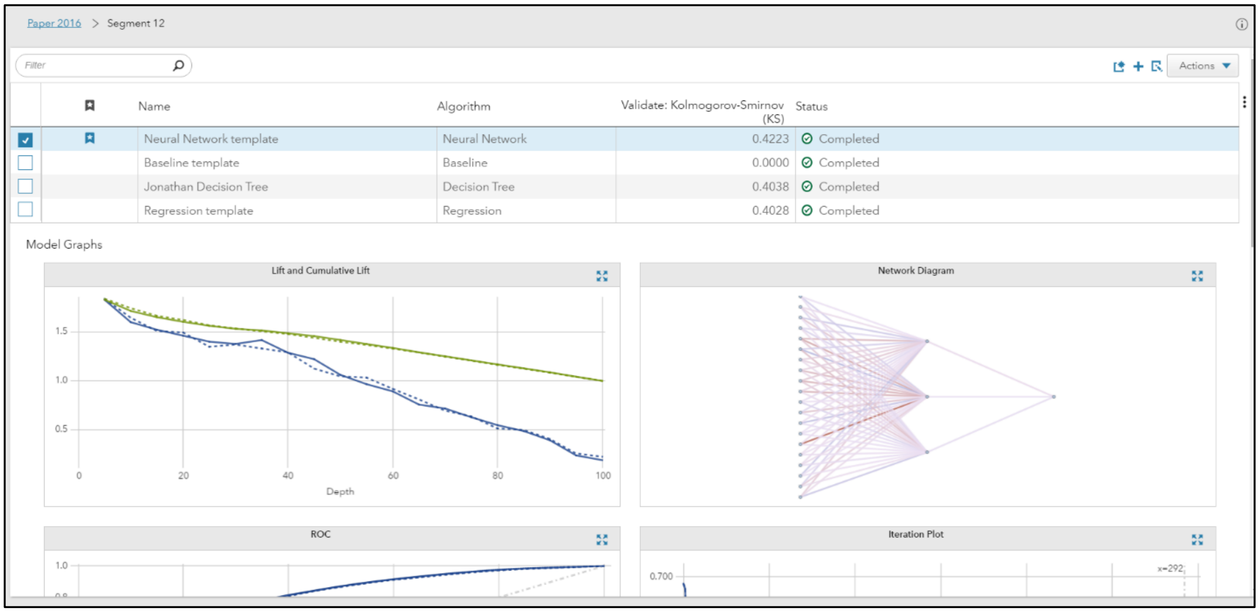

Figure 9. New York-Gold Segment Results

You have the flexibility to add new models, edit existing models, and compare models. You can override the chosen champion with any model in the segment. For example, your department might require you to have easy-to-understand results, so you might select decision tree as the champion.

To compare models, highlight the models of interest. In Figure 10, all models in the segment are compared.

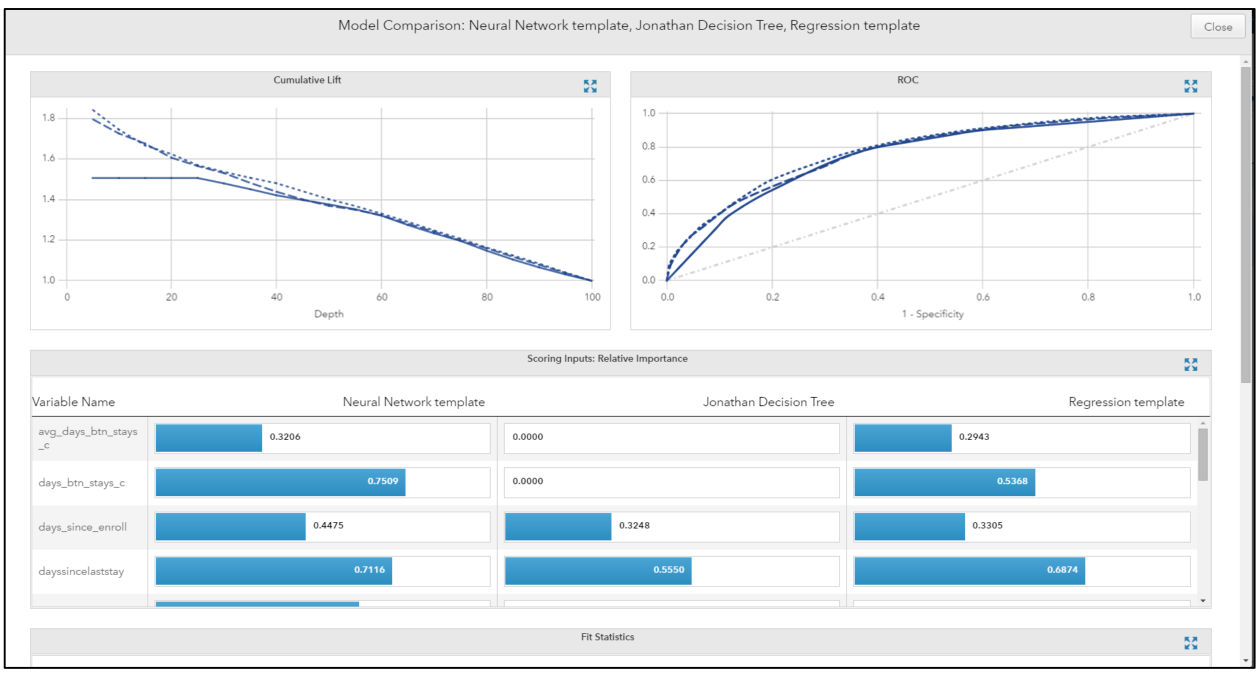

Figure 10. Model Comparison within Segment

The blue bars represent scaled variable importance within models and across models. The wider the blue bar is, the more significant the variable is within the model, relatively speaking. The variable days_btn_stays_c is more important than days_since_enroll, but not significant in the Jonathan Decision Tree.

Register Models to SAS® Model Manager

When you are satisfied that you have addressed your models appropriately, it is time to register them within the model management area of SAS® Decision Manager. You have the option to download the score code directly so that you can perform local ad hoc scoring. For automated model deployment, you can register your Factory Miner models to SAS® Model Manager, which resides within SAS® Decision Manager.

Exit the segment results screen, and then exit the project. As shown in Figure 11, to register all models, select **All models**, and click **Register**.

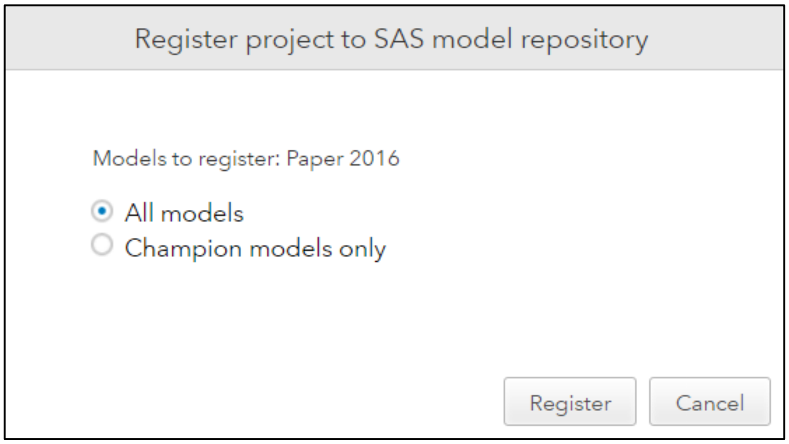

Figure 11. Register All Models to SAS® Model Manager

All models across all segments will be registered, including the Champion and any other challenger models developed using the SAS® Factory Miner templates.

All score code and required variable metadata is sent automatically to SAS® Model Manager and created within the portfolio, the model projects, and the models. Once processing is complete, a new portfolio is created in SAS® Model Manager, using the SAS® Factory Miner project name as a prefix.

SAS® DECISION MANAGER

Now that the models have been created and registered in SAS® Model Manager, we can demonstrate how, in conjunction with SAS® Decision Manager, to make decisions about handling customers from specific segments and manage those customers with the highest probability of attrition. We can also handle each segment with business rules that are customized for each segment.

So, now that the models have been registered from SAS® Factory Miner into SAS® Decision Manager, users can take advantage of SAS® Model Manager's capabilities to monitor model performance, as well as publish models to different scoring execution environments.

SAS® Model Manager

Monitoring model performance is critical to ensuring that our model is performing optimally, given the characteristics of our data. Within SAS® Model Manager, you can quickly create model performance dashboards that support automated model monitoring, dashboard updates, and the ability to automatically retrain models that have slipped below acceptable performance thresholds.

Model Portfolios

The organization and structure of content registered from SAS® Factory Miner include a hierarchy of SAS® Model Manager portfolios, model projects, and analytical models:

- Portfolios: contain the Projects registered from SAS® Factory Miner
- Projects: contains projects, one for each segment modeled in SAS® Factory Miner

- Models: contains the analytical models for each segment, one model for each analytical model created

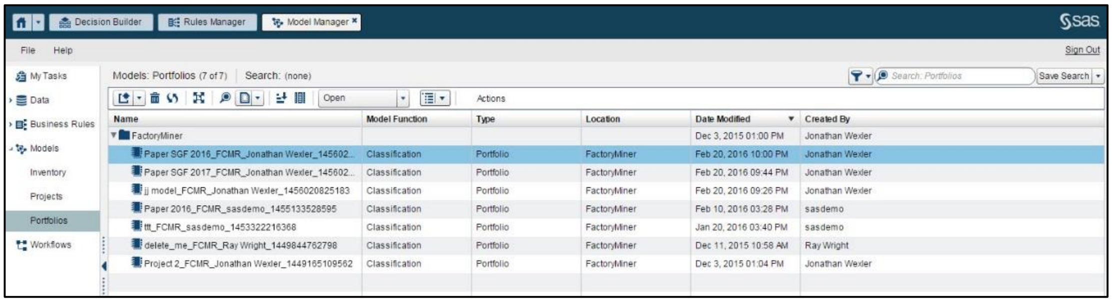

Figure 12. Portfolios Created within SAS® Factory Miner

You can drill into the portfolio to view the projects, or segments modeled in SAS® Factory Miner. Performance reports can be created easily for each Project within the Portfolio to allow you to monitor only those projects you're most interested in for reporting.

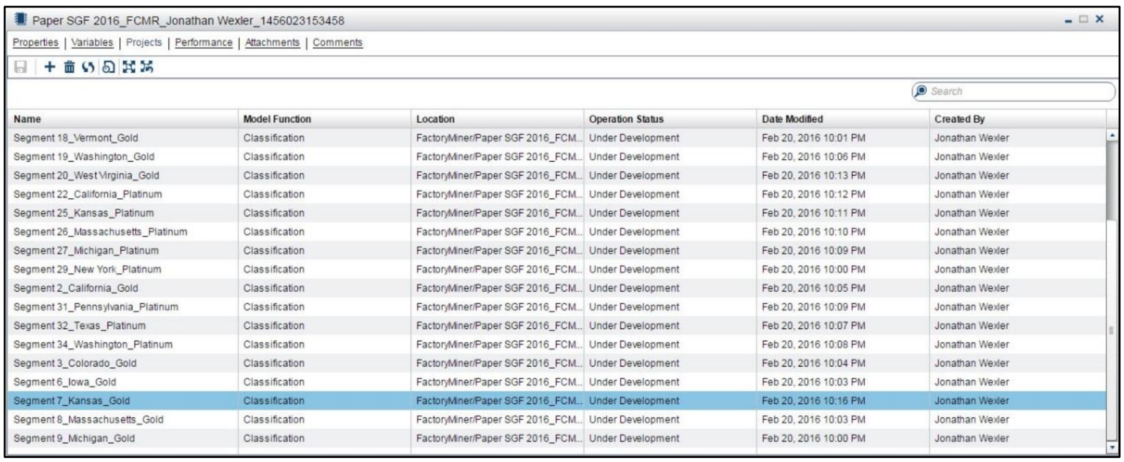

Figure 13. Segments as Projects within a Portfolio

Model Project Performance Monitoring

SAS® Model Manager enables automated monitoring and retraining when model performance degrades as data changes over time. The ways that models can be monitored fall into three main categories:

- Summary Results: The Summary results summarize the number of models, the number of versions, the number of scoring tests, and the number of reports.
- Data Composition Reports: Data composition reports include three types of reports for model project performance monitoring. The Variable Distribution report shows you the distributions for a variable in one or more time periods, which enables you to see the differences and changes over time. The Characteristic and Stability reports detect and quantify shifts in the distribution of variable values that occur in input data and scored output data over time.
- Model Monitoring Reports: The model monitoring reports are a collection of performance assessment reports that evaluate the predicted and actual target values. The model monitoring reports create several charts, including Lift, Gini - ROC (Receiver Operating Characteristic), Gini – Trend, KS, and MSE (Mean Squared Error) for prediction models.

The performance reports can be scheduled to run periodically so that you don't have to manually run each report separately. Specific model performance measurement thresholds can be specified, and when they are exceeded, warnings and alerts are delivered to users as notifications or for model retraining automation. Examples of these metrics include performance indices, output deviation indices, or model assessment metrics such as Lift decay.

Once the model is retrained and performing as expected, the model can be published to various targets or included in a Decision Flow using the SAS® Decision Manager Decision Builder.

Model Publishing Using SAS® Model Manager

The first approach to model publishing includes the ability to publish a model to a specific scoring execution environment, such as SAS® Metadata, an in-database target such as Teradata, or to an in-Hadoop target such as Hortonworks or Cloudera.

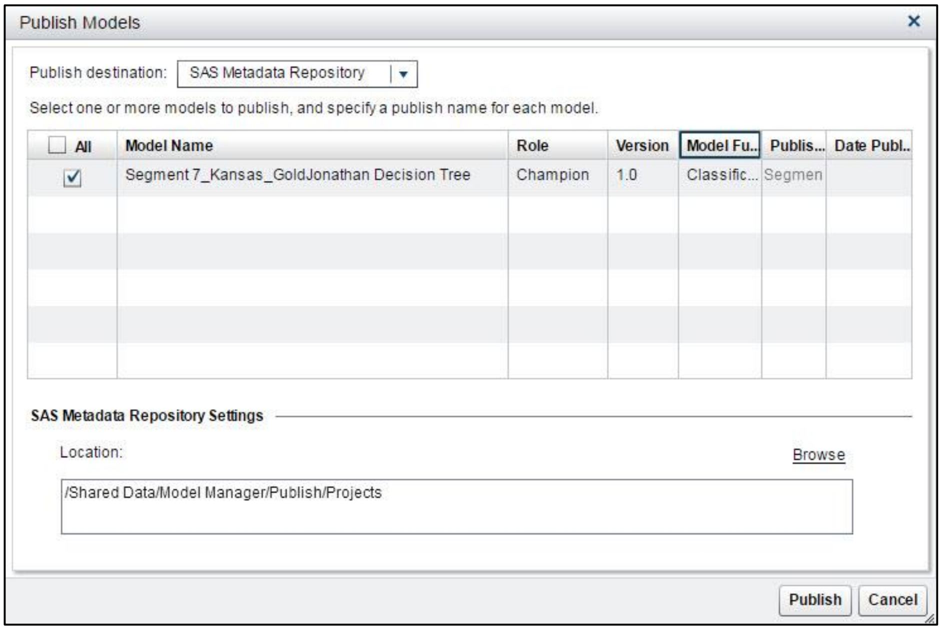

Figure 14. SAS® Model Manager Publish to SAS® Metadata

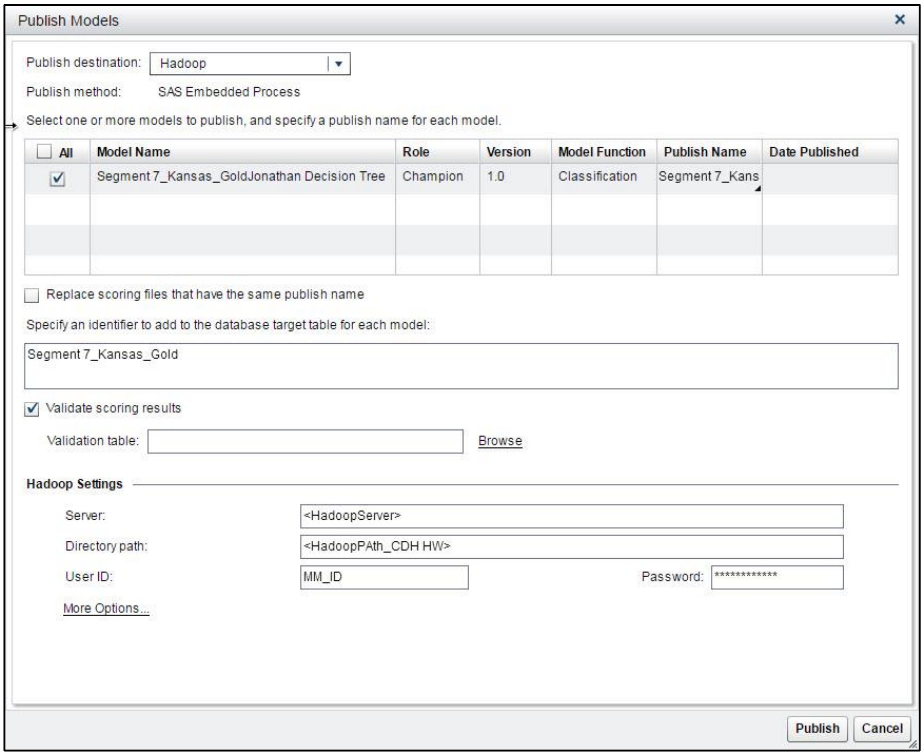

Figure 15. SAS® Model Manager Publish to Hadoop

The published models can now be discovered and data scored using these models using technologies such as SAS® Data Integration Studio for batch processing. Also, through the use of technologies such as SAS® Scoring Accelerator, model score code doesn't need to be rewritten or revalidated to score in-database or in-Hadoop, thereby simplifying and speeding the process of publishing models quickly for production scoring. Scoring specific to a segment can be accomplished across these various execution targets.

When time to market matters, you want to be able to move your models quickly into production where they can deliver results faster to improve decision making in your organization. SAS® Decision Manager and its Decision Builder takes the ability to publish models quickly one step further for better decision making that can quickly combine models and rules together into a single, easily deployed decision flow.

SAS® DECISION MANAGER AND DECISION BUILDER

Decisions within every organization are powered by powerful analytical models providing analytical precision. Using SAS® Decision Manager's Decision Builder you can easily and quickly assemble analytical models into decision flows that can be deployed for batch or web service processing within your business processes.

Decision Flow Building

We don't want our decisions to be one size fits all. Instead, we want to take actions that are specific to the segments we've identified and for which we've created models. To do that, we can use the tools available to us through SAS® Decision Manager to construct decision flows and publish them quickly as web services.

The Decision Builder user interface allows you to access the Portfolios, model projects, and models for integration into the decision flow. The user interface supports assemble models and business rules together along with conditional logic that controls exactly under what conditions your models will be executed. You can select **Add Model** to navigate to the same library to select models to include in the decision flow.

Figure 16. SAS® Decision Builder Design Interface

The model inventory displays the list of portfolios created by Factory Miner, to simplify the ability to find and integrate models.

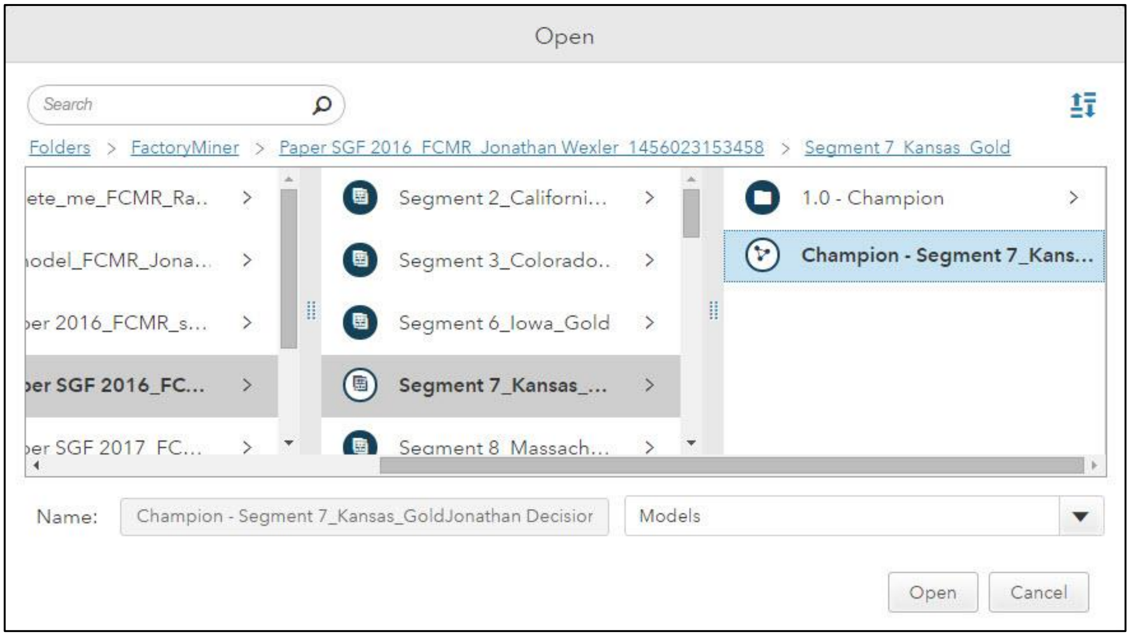

Figure 17. SAS® Decision Builder Model Inventory

The Decision Builder is building a set of variables that can be used in the decision flow develop conditional logic. In this case, we want to score and generate results based on the **member_type** and **state** in which the customer lives to decide the best action to take under these conditions. You can select a specific model, Champion or Challenger, for a specific segment (in this case, the Gold customers within the state of Kansas).

The business rules address conditions and factors that are typically outside of the model and can be integrated into the decision flow using an approach similar to the one used for integrating models. In this case, we want to consider providing an offer to customers that have a long number of days between stays (days_btn_stays_c) as well as the number of nights they've stayed as a business user (nbr_of_nights_business), and evaluate these factors along with the scores produced by the model. For this example, we want to consider customers that were considered a lower probability for an offer and apply these business rules to override the lower offer and give them a higher-value offer.

The “Rule_Flow_Loyalty” rule has been inserted at the end of the decision flow, and it defines the action we want to take for each customer (that is, define one of three options for customers and offer either a platinum, gold, or copper option based on their probability of staying with us as a loyal customer).

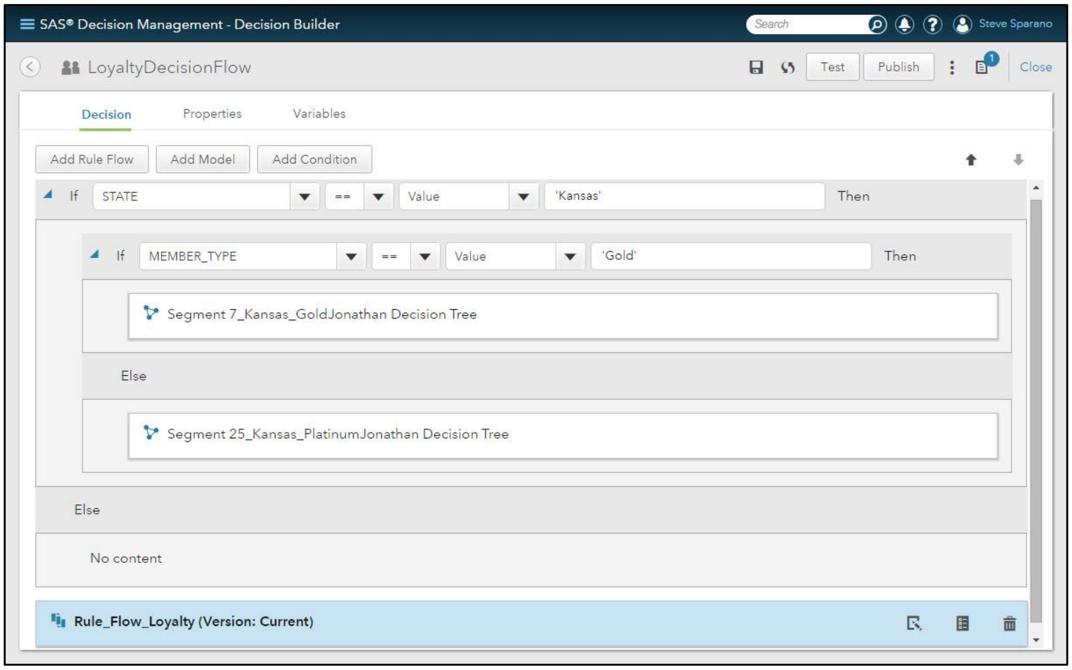

Figure 18. SAS® Decision Builder Decision Flow with Business Rules

We are able to insert specific segment models into the decision flow. This provides very granular control over the invocation of these models to produce outputs required. The conditional logic provides the control we need to execute the models and rules when they’re needed.

Publish Decision Flow as Scoring Service

Before publishing the decision flow, the Decision Builder supports testing the decision flow with data to ensure expected results match actual results. Setting up the test scenario is done by clicking the Test button in the upper right corner, as shown in Figure 18, and then selecting an input test data source. Once selected you can run the test and inspect the results, as shown below in Figure 19.

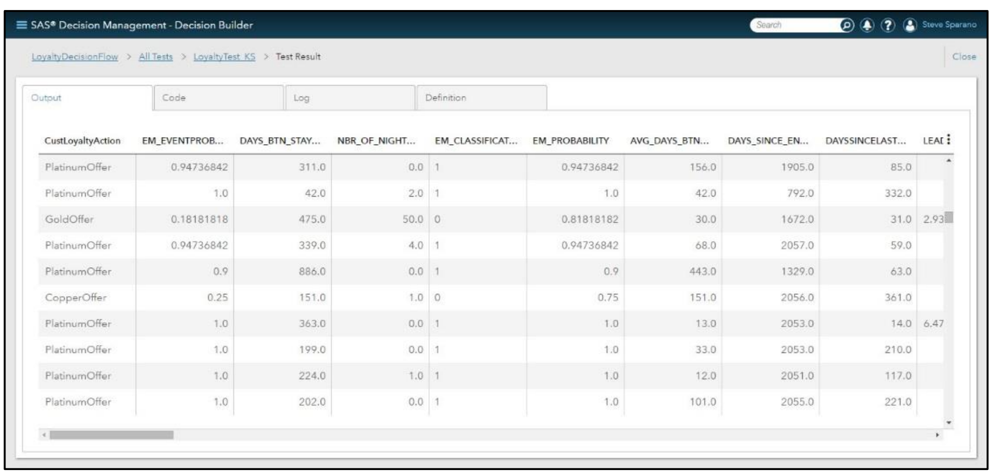

Figure 19. SAS® Decision Builder Decision Flow Test Results

During test setup, the variables required by the analytical model are required to be mapped to the table columns from the selected test table. When the names of the variables and the table columns match, the variables are auto-mapped for you. You can select the table column on the right side when they don't match to ensure the models' required input are provided during the test.

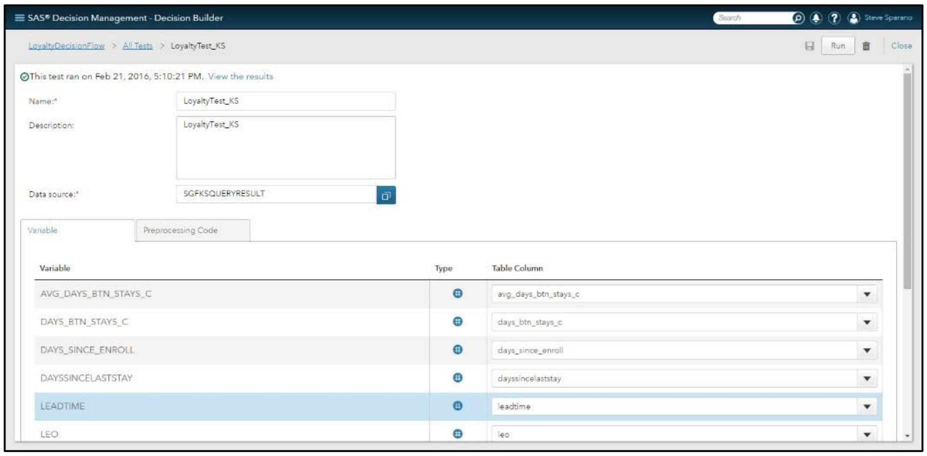

Figure 20. SAS® Decision Builder Decision Flow Test Setup

Once testing is complete, you can use the **Publish** button in the upper right corner, shown in Figure 18, to publish the decision flow to the SAS® Micro Analytic Service for creation of a scoring web service or to SAS® Metadata to enable integration with batch processes using technologies such as SAS® Data Integration Studio.

Figure 21. SAS® Decision Builder Decision Flow Publishing

At publish time, version control is provided for greater control over changes made to the decision flow. You can choose to either overwrite the latest published module (that is, version) or create a new module, thereby leaving the previously published module unaffected.

You can view the publish history to see a history of publish actions and the individuals that made these decisions.

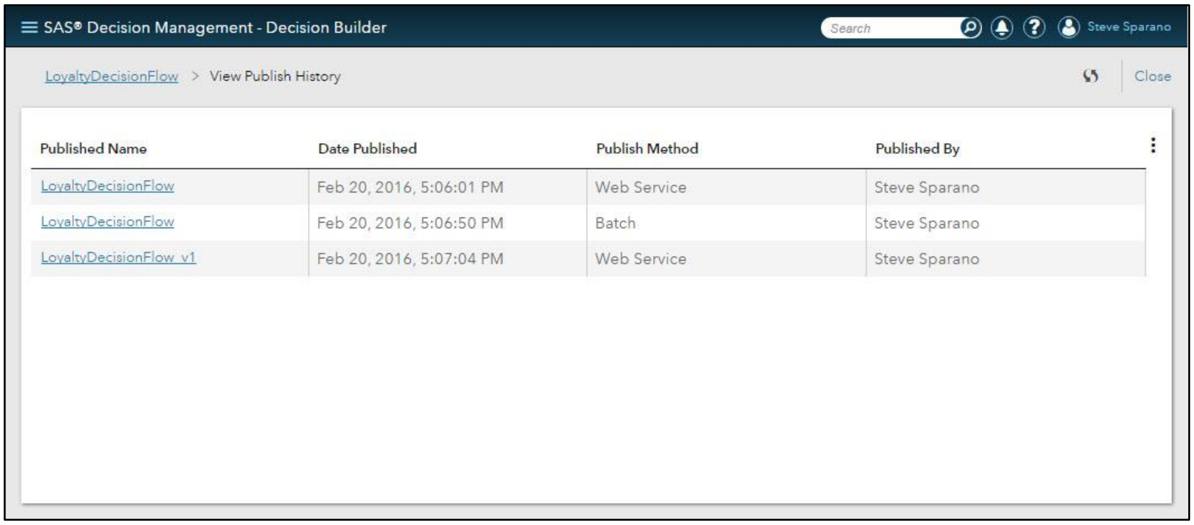

Figure 22. SAS® Decision Builder Decision Flow Publish History

In a real-time application setting, these decision flows are now available for the Loyalty applications to invoke as REST services. Applications such as call center applications will invoke these REST scoring services to deliver precise decisions about how to proceed and which actions to take for specific customer segments to ensure consistent and effective results. After all, it is within these systems and the business processes they support that operational decisions are made on a large scale.

The systems that create the analytics and related downstream decisions need to scale along with and integrate into your business processes. SAS® Factory Miner and SAS® Decision Manager provide the ability to scale as your business scales using automated and simplified deployment approaches that integrate with your business using real-time scoring services or batch processes.

CONCLUSION

Companies continue to look to analytics as a primary lever to manage their businesses. The more time you spend hand-crafting a few models and exchanging via manual handoffs, the less time you spend making an impact on your business. Time to value is critical for successful analytics. The best model means nothing unless it is an integrated part of your business.

SAS is committed to providing *modern* analytical solutions that are intelligent, scalable, and automated. Deployment of analytical models should be stateless. SAS® Factory Miner provides automated model building and collaborative intelligence to precisely target your segmented business problems. The system scales automatically to the size of your problem. Enable your data scientists to be the superheroes of your organization. Business Analysts can use their domain knowledge to enhance models while taking advantage of collaboratively built model templates.

No matter the destination, models should be deployable where the *decision* is made. SAS® Decision Manager automates and *optimizes* the process to turn models into *action.* Whether your department wants rigorous oversight of the modeling workflow, or your company wants to push models real-time, SAS® Decision Manager is your analytical center of excellence.

ACKNOWLEDGMENTS

The authors express sincere gratitude to the SAS® Factory Miner and SAS® Decision Manager developers, testers, and also to our customers.

RECOMMENDED READING

- SAS® Factory Miner Fact Sheet
  http://www.sas.com/content/dam/SAS/en_us/doc/factsheet/sas-factory-miner-107815.pdf

- Machine Learning With SAS® Enterprise Miner™
  http://www.sas.com/en_us/whitepapers/machine-learning-with-sas-enterprise-miner-107521.html
- SAS® Decision Manager Fact Sheet
  http://www.sas.com/content/dam/SAS/en_us/doc/factsheet/sas-decision-manager-106488.pdf

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors at:

Jonathan Wexler
SAS Institute Inc.
100 SAS Campus Drive
Cary, NC 27513
Email: jonathan.wexler@sas.com

Steve Sparano
SAS Institute Inc.
100 SAS Campus Drive
Cary, NC 27513
Email: steve.sparano@sas.com