



DECISION
MANAGEMENT
SOLUTIONS

James Taylor
CEO

Three steps to put Predictive Analytics to Work

The most powerful examples of analytic success use Decision Management to deploy analytic insight in day to day operations helping organizations make more profitable operational decisions.

Organizations are increasingly adopting predictive analytics, and adopting these predictive analytics more broadly. Many are now using dozens or even thousands of predictive analytic models. These models are increasingly used in real-time decision-making and in operational, production systems. Models are used to improve customer treatment by selecting the next best action to develop a customer, to make loan or credit pricing decisions that reflect the future risk of a transaction, to predict the likelihood of equipment failure to drive proactive maintenance decisions, or to detect potentially fraudulent transactions so they can be routed out of the system before they hit the bottom line. Examples like these deliver high ROI from analytics.

However, many analytic teams rely on approaches and tools that will not scale to this level of adoption. These teams need a repeatable, effective, efficient process for creating and deploying predictive analytic models into production. They must *operationalize* analytics.

Operationalizing analytics has three elements:

- ▶ A collaborative environment and shared framework for problem definition to ensure that the analytics created are solving the right problem.
- ▶ A repeatable, industrial-scale process for developing the dozens or even thousands of predictive analytic models needed.
- ▶ A reliable architecture for deploying and managing predictive analytic models in production systems.

The solution to operationalizing analytics involves the effective combination of a Decision Management approach with a robust, modern analytic technology platform. Such a combination focuses analytics on the right problems and effectively integrates analytical results directly into operational systems for faster and more profitable decisions.

This paper discusses both how to use a focus on decisions to ensure the right problem gets solved and what such an analytic technology platform looks like.

CONTENTS

- Solving the Right Problem
- Creating an Industrial-Scale Process
- A Reliable Deployment Architecture
- Conclusion



Solving the Right Problem

The first step in most analytic projects is developing a model plan. A key challenge in developing an effective model plan is ensuring the business problem to be solved is well understood. How will I measure success? What data is applicable? Who will use this model and what decisions will it influence? Where will I need to deploy this model? An effective analytics team works in close collaboration with their business partners to answer these key questions and stays in synch through the model development and validation process. Working closely with the IT team also matters and improves data access while easing the path to deployment. However, there is often a gap in understanding, tools and collaboration environment between the business, analytic and IT teams.

Today, however, most organizations have ad-hoc processes for business problem definition in their analytic projects. As a result, projects lack a common vocabulary and there are few if any tools for collaboration with the business and IT teams. Gaps in business understanding lead to analytic models that solve the wrong problem or solve the right problem the wrong way. A propensity model that relies on data not available in the marketing system is of little use or a model designed to influence a decision but which must be made exactly as specified in a regulation needs to include appropriate rules. It can take months for predictive models to be deployed and some models never get deployed at all. Good analytic models too often just sit on the shelf and analytic resources are wasted. Other models get deployed but are not monitored and this can result in models that degrade to the point where they do more harm than good.

YouSee is Denmark's leading provider of cable TV and broadband services and is a division of the TDC Group, Denmark. YouSee faced two challenges. First, customers of TV and broadband products do not tend to be very loyal and customer churn is a constant problem. Second, multi-product customers are more profitable and more loyal but many YouSee customers have only a single product. A project to bring predictive analytics to bear on these challenges was launched.

YouSee was clear what decision it wanted to improve—the decision that a call center agent makes about a cross-sell or retention offer when talking to a customer. This focus allowed them to identify the data that would drive their models and create two initial models that they knew would improve this decision—the likelihood of a successful cross-sell of cable TV services to a broadband subscriber in the next 90 days and the likelihood that a broadband subscriber will churn in the next 90 days.

The first step in operationalizing analytics is developing an understanding of the problem, an understanding that will allow the problem to be solved. Decision

Management is a proven approach that creates a shared framework and collaboration environment for the business, IT and the analytics teams. Within this framework they can come to an understanding of the problem by identifying and prioritizing the operational decisions that drive the organization's success.

Decision Management links operational decisions to the business drivers and performance measures that have the most impact on the business. Modeling the operational decisions to be analytically improved clarifies and focuses the problem statements that drive analytic projects. This focus ensures projects show a strong and lasting ROI and effectively apply constrained analytic resources. Because it is a business oriented approach, Decision Management creates productive interaction between the analytics and business team. Analytic teams report it is much more interactive and effective than current processes to gather requirements.

Decision Management also identifies critical operational decisions currently buried in business processes and IT systems. Describing and visualizing previously embedded operational decisions as separate, discrete activities helps business stakeholders understand and take ownership of how decisions are currently being made. Linkage to the business process and existing IT systems clarifies how and where the resulting analytics will need to be deployed, helping guarantee a successful deployment once analytic models are developed.

“It really helps our auditors work smarter. You don't eliminate the person; you help them pick which claims we should look at, and for what reason.”

Gary Hewes, General Manager of Service Support Operations, GE C&I

Creating an Industrial-Scale Process

The next step in operationalizing analytics is moving to an industrial process for building analytic models. This means moving away from individual scripting environments where every task is performed by hand, reuse is limited, only small numbers of expert analytic practitioners are involved and these practitioners do not follow a standard process. Such an approach can and does produce high quality models but it cannot scale to enable an organization to become broadly analytic in its decision-making.

An industrialized process incorporates five key characteristics:

- ▶ A systematic approach to data management.
- ▶ A predictive analytics workbench environment that supports reuse, automation and repeatability.
- ▶ Engagement of less technical users.
- ▶ High performance analytic architecture for fast model turn around.
- ▶ Ongoing management and monitoring of models.

A systematic approach to data management

The increasing volume, variety and velocity of data are making data management more complex—an increasing number of analytic problems are Big Data problems. Organizations need to be able to use transactional data, transactional event streams, unstructured or semi-structured data and more to understand what customers are doing, for instance, or what constitutes an unacceptable risk. A systematic approach to information management that encompasses data quality, data integration and data management ensures access to data is standardized and efficient. Adding Master Data Management ensures well defined metadata and ensures modelers are working with standard sets of data about customers, products, etc. Data definitions are shared and analytical datasets are generated in a repeatable and increasingly automated way, speeding the production of analytic models and improving their consistency.

A flexible platform for data management also allows the right kinds of data to be brought to bear each time, whether that data is a data warehouse, flat file or NoSQL store. SAS / Access® Interface to Hadoop, for instance, allows access to and from Hadoop using the SAS language and tools as well as delivering Scoring Accelerator for Hadoop. SAS Metadata allows effective management of metadata across multiple data sources including Hadoop. Finally all these elements must be managed through a common governance layer.

A predictive analytics workbench environment

This more systematic approach to information management feeds into a workbench environment for defining the modeling flow. Products such as SAS® Enterprise Miner™ allow reusable and repeatable model creation flow diagrams to be developed and shared through a repository. These workflows streamline and standardize how analytic modeling is performed. In-database mining capabilities integrated with these workflows can push data preparation, transformation and even modeling algorithms into an organization's data infrastructure, improving throughput by reducing data movement. In-memory and other high performance analytic capabilities as well as intelligent automation of modeling activities can be applied as appropriate in the steps defined in the workflows.

Engagement of less technical users

Modeling capacity is a critical challenge for any organization that is broadly adopting analytics. Historically, analytic modeling has required specialist resources. These workbenches, however, can often be wrapped with an interface suitable for less technical users. Features like those in SAS® Rapid Predictive Modeler allow less technical users to build and execute models that take advantage of underlying automation capabilities to produce large numbers of “good-enough” models quickly. Working collaboratively with an analytic team, these users can produce first cut and less core analytic models, participate more fully in reviews of models and allow the analytic team to focus on high-value, high complexity problems.

High performance analytic architecture for fast model turnaround

Organizations seeking to broaden their use of predictive analytics also face challenges with the rapidly growing data volume and complexity and slow time to results because underlying hardware infrastructure is maxed out. A high performance analytic infrastructure scales out while delivering high availability, workload management and scheduling. Products like SAS® Grid Manager create a distributed grid environment that provides parallel job execution across multiple servers with shared physical storage. Provided algorithms have been written to take advantage of this infrastructure, dramatic performance improvements are possible. Products like SAS® Data Integration and SAS® Enterprise Miner™ are automatically configured to support this kind of parallel processing in a grid computing environment. Other SAS solutions, like SAS® Risk Dimensions® and SAS® Forecasting, can be set up to automatically submit SAS jobs to a grid of shared computing resources. For everything else, like Base SAS code, a few lines of code “grid-enable” it.

One example is a portfolio risk analysis scoring on 400,000 loans reduced from 3 hours to 10 minutes and another where processing time for customer behavior models (segmentation, propensity and retention models) was reduced from 11 hours to 10 seconds.

Combined with powerful visualization tools like SAS® Visual Analytics, in-database analytics, and in-memory analytics, these techniques can make new problems solvable and increase the capacity of an analytic team by reducing turnaround time for each iteration of model creation.

Chico's, based in Fort Myers, FL, operates more than 1,000 boutiques throughout the US, US Virgin Islands and Puerto Rico under the Chico's, White House | Black Market and Soma brands. The company also markets to consumers through catalog and online channels. Chico's segments its catalog mailings and differentiates promotion efforts for maximum impact using SAS OnDemand: Marketing Automation.

This cloud-based predictive analytics solution allows Chico's to send customers with a propensity to buy items as soon as they are released a full-size, full-price catalog and mailings highlighting new merchandise. Discount shoppers get slimmer liquidation catalogs and sale fliers. Online customers receive e-mails geared to their buying habits.

Ongoing management and monitoring of models

Finally, predictive analytic workbenches also support the ongoing management and monitoring of models once they are in use. Capabilities such as those in SAS® Model Manager™ allow an analytic team to set up automated monitoring of models to see when they need to be re-tuned or even completely re-built. These capabilities also help track the performance of models to confirm their predictive power and behavior. This workflow too can be defined and managed, bringing even those not using the predictive analytic workbench, such as database administrators and integration specialists, into the process.

It's easy to think this is all about technology but it's not. Effectively operationalizing analytics requires the right technology, the right people and new processes. Only focusing on the technology will not be enough.

A Reliable Deployment Architecture

Today, the results of the predictive model development process are often captured in a document describing the model, in a report or dashboard, or in code generated by the modeling system. For strategic decisions, a report is often sufficient as the person who needs the insight can review the analysis and make their decision. But to harness the full value of predictive analytics, analytic insights need to be deployed in operational systems.

A model may be hand coded into a system for operational deployment. Once upon a time companies would re-type information to move it from one IT system to another—rekeying it because the systems were not integrated. This kind of inefficiency may be a thing of the past but many organizations using analytics are doing something very similar with their analytical model results. In the process they are creating huge manual delays, increasing the risk of errors and reducing the accuracy of their models.

Furthermore, model implementation and updates typically require multiple IT projects each of which take time and cost money. Even when the analytic tool produces code, this code and associated data become quickly embedded and entangled with other system's code, making subsequent course corrections difficult and costly.

Deploying models into operational systems is central to generating significant business value. But organizations regularly find that 50-60% of models aren't deployed and that deploying the remainder takes 30, 90, even 120 days or more. Businesses cannot bear the opportunity cost of not deploying and using these analytic insights and organizations cannot afford this level of inefficiency in their use of constrained analytic resources.

YouSee's decision focus ensured that it was clear all along that simply displaying predictive probabilities in the call center application would be unsatisfactory. Integration of the predictive analytic models into the call center application included developing business rules that used the scores and other customer data to generate dynamic, customer-specific scripts within the call center system. These scripts have allowed individual call center agents to achieve a significant increase in success rate on the cross-sell suggestions.

With an industrial-scale model development process addressing the right business problems, Decision Management takes predictive analytics to the next step by enabling a reliable architecture for deploying predictive analytic models into operational systems. Decision Management successes across industries demonstrate the increasing value of predictive analytics when applied to the quality and effectiveness of operational decisions. Making faster, better and more consistent

decisions requires the integration of decision logic and analytics in operational systems.

Operational decisions are highly repeatable decisions about a single customer or transaction. Because they are made over and over again, these decisions create the data needed for effective predictive analytics. Because operational decisions are made by front line staff or completely automated systems they don't lend themselves to analytic approaches that rely on the end user such as visualization and query tools. To influence these decisions using predictive analytics, executable analytic models are integrated into operational systems, driving better decisions in real-time in high-volume, transactional environments.

“SAS is a hard-wired part of our operations”

Greg Hayden , Head of Commercial Analysis & Information Systems, Frontline

Decision Management ensures that when a model is built and validated all the teams involved, the business, the analytics and the IT teams, know where and how the model will be deployed. The teams can now focus on the point at which operational decisions are being made and on the systems used to make those decisions.

Modeling operational decisions defines the logic of how a particular operational decision should be made given the system's understanding of the current situation. Business Rules Management Systems such as the SAS® Rules Manager, ensure the decision-making logic is both managed and executed transparently, so it is clear how each specific decision was made. This decision making logic can then take advantage of predictive analytic models deployed using in-database scoring infrastructure or directly into operational environments.

In-database analytic deployment

One way to make predictive analytic models available in operations is to use in-database scoring infrastructure such as the SAS® Scoring Accelerators. These take analytic models and push them directly into the core of your data stores, including Hadoop. Once deployed, the models are available as a function or as SQL and can be included in views or stored procedures. This allows operational systems direct access to the result of the model while ensuring that this is calculated live, when requested, and not based on a potentially out of date batch run. The business rules for a decision can then access the predictive analytic result like any other piece of data.

Direct analytic deployment

Predictive analytic models can also be deployed directly into operational environments using Decision Services. Decision Services combine the decision logic that represents the expertise, regulations and policies that drive a business decision

with predictive analytic model results. Decision Services are the implementation of a decision—how other systems will find out what the decision is. A Decision Service also makes the decision reusable and widely available. Decision Services are essentially business services in a Service Oriented Architecture (SOA) that deliver an answer to a specific question. These services generally do not update information—they just answer questions. Because they don't make any permanent changes, they can be used to answer questions whenever they come up without worrying about potential side effects.

Some predictive analytic techniques such as decision trees and association rules can be expressed directly as decision logic or business rules that execute in the decision service. The increasing adoption of the Predictive Model Markup Language (PMML), an XML standard for defining predictive analytic models, means that analytic workbenches (including SAS® Enterprise Miner™) can read and produce PMML and then deploy into an execution environment. Many business rules and business process management systems including SAS® Decision Manager support PMML. This allows them to load the definition of a model directly and then execute it when the rules or process need to know the prediction.

In addition many packaged applications such as those for fraud detection, credit risk management, or customer intelligence have been designed to allow new models to be rapidly integrated into operational systems. As more packaged applications tackle decision-making this capability will only become more widespread, giving analytic teams more deployment options and further reducing the barriers to using predictive analytics in operations. SAS® Customer Intelligence and SAS® Fraud Detection are two examples.

Staying predictive and effective

A reliable deployment architecture is only as effective as its ability to stay in synch with the changing business environment. In addition to model monitoring and tuning, the application of performance management techniques and technologies to the monitoring of decisions is critical. With Decision Management, the business now understands how specific decisions create value. These decisions are linked to the business and individual performance metrics being tracked. To continuously improve business performance, a Decision Management approach monitors decision performance, throughput and basic statistics. How many decisions are made to approve, reject or refer is a measure of decision effectiveness. Too many referrals will increase the burden on staff doing manual reviews. Too many rejections, thanks to false positives for instance, will impact customer service or sales. Similarly, decisions that take too long or that cost too much (because they use data that must be purchased, for instance) may have a negative overall impact. Tracking and

reporting on this information will help the business owners understand and thus manage their decisions more effectively.

This kind of decision performance analysis is easily combined with the model monitoring capabilities increasingly found in model deployment infrastructure like SAS® Model Manager. This enables drops in decision performance to be tied to changes in model behavior thereby driving the right kind of model updates. Rapid deployment of the changed model, thanks to strong understanding of the decision-making context and modern deployment infrastructure, ensures that models stay current, predictive and effective.

Real-time

Regardless of how the predictive analytics are deployed into the Decision Service, they should be used in real-time. That is the scores and other results from the predictive analytic models should be calculated as and when needed to make a decision. This ensures that the most up to date data is always used. Recognizing that modern operational systems are real-time is critical for analytic teams. Too many predictive analytic models are currently deployed using batch scoring, resulting in models calculated without the most current data.

The GE C&I Consumer Home Services Division estimates it saved \$5.1 million in the first year of using SAS to detect suspect claims. GE Consumer & Industrial (C&I) relies on thousands of service providers to handle more than 1 million service claims in the appliance group each year. GE's Consumer Home Services Division used to audit only a portion of claims that fell within certain parameters. By operationalizing analytics that predict the likelihood of fraud, GE now automatically audits 100 percent of the claims for suspicious activity. The predictive analytic models embedded in the operational environment evaluate facts about the services performed including the product, customer, complaint or work completed to predict how likely it is that this claim is fraudulent. The audit team can then focus their efforts on the highest risk claims.

Conclusion

Predictive analytics can deliver insight, predictions, and resolve business uncertainty into profitable probabilities. To take advantage of this insight and make better, more profitable decisions, organizations are deploying predictive analytics in their operational systems. Decision Management combined with a robust and modern technology platform delivers:

- ▶ A collaborative environment and shared framework for problem definition to ensure the analytics is solving the right problem
- ▶ A repeatable, industrial-scale process for developing the dozens or hundreds or more of predictive analytic models that will be needed, and
- ▶ A reliable architecture for deploying predictive analytic models into production systems.

Decision Management improves operational decisions by deploying predictive analytics in Decision Services that support many operational processes and systems. Decision Management focuses on decision effectiveness and more effective decisions reduce fraud, improve customer interactions and spot opportunity.

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration. Other brand and product names are trademarks of their respective companies.
101 | H_S1G | JI. 01 14

Contact Us

If you have any questions about Decision Management Solutions or would like to discuss engaging us we would love to hear from you. Email works best but feel free to use any of the methods below.

Email : info@decisionmanagementsolutions.com

Phone : +1 650 400-3029

Fax : +1 650 352-9247

