Paper 125-2011

Using Real-Time Analytics to Manage Next Best Action

Chris Toth, SAS Institute Inc., Boston, Massachusetts

ABSTRACT

How can you continually increase the value of customer interactions? Typical campaign management and batch processes look only backward. They do not account for the real-time context of the customer interaction. For example—the reason for a call or visit to your Web site, the demeanor of the customer, the product interest that the customer has—these factors and characteristics are dynamic and evolving. This paper discusses how SAS® Real-Time Decision Manager helps you make your offers relevant to the interaction and behavior, which, in turn, increases customer satisfaction and improves offer response rates.

INTRODUCTION

If behavioral data is the key to understanding people, then analytics is the key to developing insight. Furthermore, real-time analytics is the key to delivering immediate, actionable customer insight. The gaming industry is famous for collecting enormous amounts of information about its patrons. Analytics is what successful gaming organizations use to distill, enrich, and leverage that information. Each interaction with a customer might be the last, so real-time analytics enables organizations to leverage current and past events and behavior to deliver the next best action. Delivering the next best action fulfills an organization's goals, avoids conflicting goals, increases lifetime value, and ensures a unique response to each unique situation.

Satisfaction and lifetime value are the most important metrics when determining the next best action. Other metrics contribute to satisfaction *or* lifetime value, but not to both. Satisfaction drives brand loyalty and referrals. Brand loyalty and referrals drive the timing and frequency of transactions. All of these contribute to lifetime value. Another factor that is important is a customer's social influence score. Keeping highly networked customers extremely satisfied not only raises their own value higher, but also raises the value of customers connected to them. An organization should constantly strive to improve satisfaction and increase lifetime value in every customer interaction, particularly with customers who are well connected.

With the technology available in today's society, customers demand instant gratification. A customer needs to know he is getting the best value for his need, the lowest cost, the highest quality, the best-suited product or service. When a customer is in the process of making a purchase, an organization usually has only one chance. A long-standing relationship with this customer only means that the customer comes to you first, before shopping around. You need to ensure that your offer is the most relevant based on the customer's preferences. Otherwise, the customer will take her business somewhere else.

THE RATIONALE FOR REAL-TIME ANALYTICS

CUSTOMER RELEVANCE

Customers interact with organizations through many channels, the most common being brick and mortar stores, call centers, the Web, e-mail, chat rooms, mobile devices, Facebook, and Twitter. Consistency in these channels is critical to maintain relevance with the customer without oversaturating the customer. Messaging systems need to retain what messages have been sent to avoid sending duplicates. They need to reinforce critical information, and they need to learn from responses. This knowledge helps organizations provide the next best action for the customer. Industry-leading organizations are pushing customers toward lower-cost channels, such as the Web. They promote these lower-cost channels through discounts or personalized portals that incorporate recent behavior. As a result, the interaction becomes more relevant for the customer.

Relevancy is based on insight, and analytics enables organizations to capitalize on a customer's multifaceted personal preferences, recent activity, and objective.

Personal preferences constantly evolve. Preferences can include price sensitivity, brand loyalty, and liking a certain product tier (such as budget, mid-line, or premium). Preferences can include product reliability, service plans, extended warranties, or products that are geared toward novices or experts.

Recent activity includes Web site searches, queries, and complaints to the call center. It includes feedback on Facebook or Twitter. A customer is receptive to advice or an offer that helps her make a final purchase decision based on what is most important to her. If a customer is using an interactive channel such as a store, call center, or chat session, then a representative has the golden opportunity of having a meaningful and iterative conversation. However, if a customer is using a Web browser, then there is a single chance to offer the next best action.

Industry-leading organizations are using real-time analytics to decipher customer preferences, historical data, recent behavior, and information to determine the product or service that is most relevant. As a result, the customer is more likely to “make the purchase.”

EVENTS, BEHAVIOR, AND THRESHOLDS

Events and behavior are easy to categorize with enough data. It’s the thresholds at which actions are manifested that are difficult. For example, I might have five dropped calls in a single day on my mobile phone if I am in a bad reception area. I might get frustrated and tweet about it, but this is the extent of my reaction. In contrast, my friend, Bob, has three dropped calls in one week, but these calls are important business calls. Bob becomes so irate that he calls his provider to complain. His frustration continues as he speaks to a representative who offers him 500 free SMS messages as compensation for his troubles. The problem is, Bob barely texts. Bob cancels his service and signs up with a competitor. We all experience similar positive and negative events. We remember a compassionate gesture, a two-hour online super sale, a dropped call, an Internet outage, or a defective product. Our behavior falls into typical categories—we purchase a deeply discounted product, you tolerate a network glitch, or I cancel a subscription. Behavior is shaped by cumulative satisfaction. A man is likely to stay with a barber if he consistently gets a great haircut. If he gets a bad haircut, he might comment on it, but as long as the recovery is good, and the great service continues, he stays with the barber. When he gets multiple bad haircuts, he is dissatisfied, the experience is negative, and he ends the relationship. On the flip side, a woman is typically more sensitive about her hair. One bad cut can end the relationship with the hairdresser. What makes us truly individuals is our thresholds to stimuli or the points on which we take action. At what point is a discount too good to pass up? At what point is a dropped call too much to bear?

Predictive analytics determines these thresholds. It knows our propensity to accept an offer or the risk of attrition. This knowledge, combined with the lifetime value of a customer, determine the next best action. Almost anyone would purchase an iPad at a 90% discount. Almost anyone would sign up for a two-year contract with this iPad. But, how *many* people and *which* people would sign up for a two-year contract for a 5%-off iPad? At what discount would a person sign up, and how many additional products or services would he or she purchase over the next three years? How many friends could this person influence to make the same purchase? These are questions industry-leading organizations use analytics to answer. Real-time analytics provides the next best action immediately and uses current information. Real-time analytics incorporates all contextual information and calculates the propensity a customer has to accept an offer or the risk of attrition. To do this, real-time analytics uses recent events, customer behavior, demeanor, and the purpose of the Web visit or call.

WHY REAL-TIME ANALYTICS?

If analytics is the key differentiator in successful and sophisticated organizations, then why are most of these organizations scoring their customers only monthly, and adding a transient score overlay for daily events? Organizations are just now quantifying the opportunities that are lost. It’s critical to incorporate significant events in a timely manner, and to score on demand so that decisions can be made with up-to-date information. If a credit score is updated only on a weekly basis, and a person declares bankruptcy, then that person could potentially benefit from refinancing his debts before the bankruptcy is reflected in his credit score.

Customer satisfaction can be fleeting. Organizations must be insightful about their customers, and they must tactically act on that insight in a timely manner. To determine the next best action for a customer, an organization must consider the customer’s lifetime value and the customer’s propensity to accept the offer. The organization must leverage the customer’s current satisfaction, which can directly affect the risk of attrition. Satisfaction is a vague metric, but analyzing past behavior, responses, recent events, and demeanor can quantify this vague metric and make it an invaluable asset.

WHY NOT REAL-TIME ANALYTICS?

It’s difficult to operationalize analytics. Today’s applications are typically a heterogeneous blend of legacy systems and new service-oriented architectures. Data is stored in the applications themselves, and in various databases. After real-time analytics scoring is integrated into this complex mix, it can take months from when you create the model to when you deploy it into production. IT departments are already strained with application development and maintenance, so adding analytic coding and deployment cycles could be impossible without hiring new resources. Organizations that have these resources still need three to six months to deploy the model because of the strictness of the Software Development Life Cycle (SDLC) and application release cycles. Although the benefits of real-time analytics are clear, traditional IT processes and logistics make it slow and difficult.

THE EVOLUTION OF OPERATIONAL ANALYTICS

In the beginning, organizations manually coded scoring algorithms in their programming language (C, C++, Fortran, Pascal, Java). Significant testing was required to ensure that the models were properly implemented and integrated, and that the entire system worked as expected. The result was many months between a completed analytic model and that analytic model being implemented, integrated, tested, and deployed into production. SAS introduced code generation for model output with SAS Enterprise Miner, and the development cycle was dramatically shortened. Months of implementation were cut by eliminating the manual coding and implementation of the models. However, there was still integration and testing. The result was fewer months between a completed analytic model and that analytic model being implemented, integrated, tested, and deployed into production during the application's next release cycle.

Business process automation (BPA) works in parallel with analytics to provide business efficiencies. BPA is often enabled with technologies such as business rule engines (BRE) and workflow engines. BREs enable organizations to remove hardcoded business logic from their operational systems. BREs empower business users to create and maintain their policies, practices, and procedures in a business-friendly environment. BREs deploy updates to the business logic in parallel to application releases and independent of IT schedules. This functionality enables IT to focus on infrastructure, and the organization to focus on revenue.

In late 2007, SAS introduced SAS Real-Time Decision Manager. This product reinvented the process of operationalizing analytics. Learning from the success of BREs, SAS Real-Time Decision Manager works with SAS Model Manager to seamlessly deploy analytic models from SAS Enterprise Miner via Web services. Once the initial integration is completed, a model update can be deployed without restarting or modifying any operational system. The result is *days* between a completed analytic model and that analytic model being implemented, integrated, tested, and deployed into production, independent of the application's release cycle. As a result, SAS Real-Time Decision Manager enables IT to focus on infrastructure, and the organization to focus on revenue.

Operationalized models can be reused in multiple applications and channels, which leads to consistent decision making across all customer touchpoints. For example, a credit risk score that is used by the call center system can be leveraged on the Web in the banking portal.

EMPOWERING MARKETERS, BUZZWORDS, AND REALITY

OFFER ARBITRATION

Offer arbitration is the process of prioritizing competing offers to determine the next best offer based on a set of eligibility criteria and weighting. Customer relationship management systems have used offer arbitration for years with limited success. While the concept of offer arbitration fulfilled business requirements and expectations, in practice, it rarely lived up to expectations. The eligibility criteria would work flawlessly. On the other hand, the arbitration algorithms were often simple sorting algorithms that used batch scores that were not current, and used empirical weighting schemes and rules that did not accurately reflect a customer's intentions.

The SAS approach to offer arbitration was the product of a decade of practical experience. Just as operationalizing analytics is the differentiator of a successful organization from a business process automation perspective, predictive analytics is the differentiator when it comes to relevant offer arbitration. It is true that SAS uses simple sorting algorithms for organizations that are migrating from systems that use traditional offer arbitration. However, the power of the SAS approach is real-time predictive analytics. The SAS Customer Intelligence product suite combines real-time online behavior collection and classification with real-time market-basket analysis and

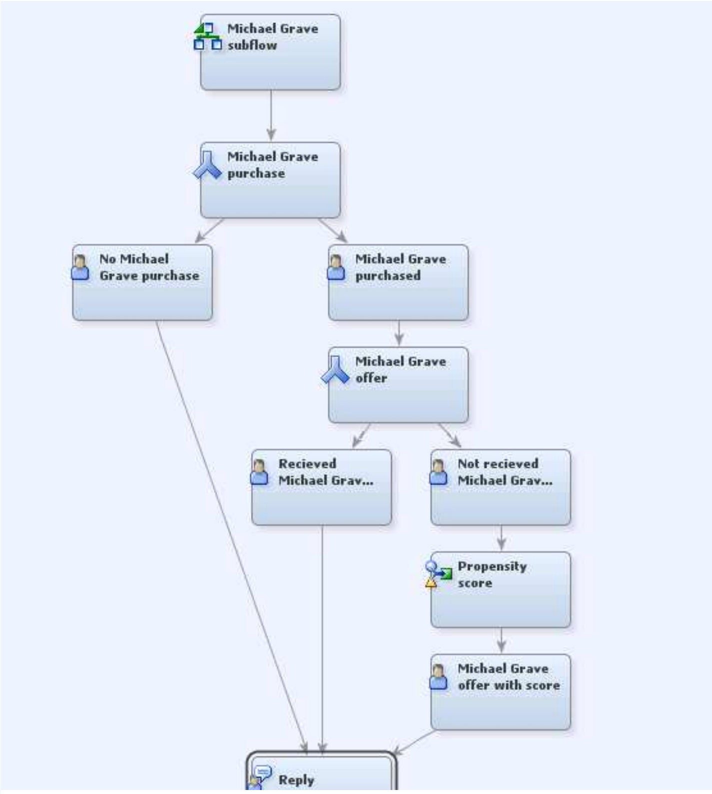

Figure 1~~2~~: Offer Arbitration within SAS Real-Time Decision Manager

offer-propensity scoring. SAS enables organizations to align with their customer at every touchpoint, during every interaction, with real-time analytics insight.

SELF-LEARNING ALGORITHMS

The promise of self-learning algorithms is a zero-maintenance analytical model that updates itself using customer responses, and that provides heavily optimized performance from a scalability perspective. These self-learning algorithms are fit for purpose using a black-box approach, and they work well for a very specific task and for the propensity to accept an offer. But, they have little applicability outside of their narrow focus for uses such as attrition risk or probability of default. Self-learning algorithms yield fairly simplistic and obvious results. For example, if a customer purchases a children's book in December, then a self-learning algorithm suggests you offer another children's book. These algorithms can be distorted by seasonal and localized trends, which leads to a self-fulfilling prophecy—substantial holiday season children's book sales. Self-learning algorithms are anonymous by nature, and they do not incorporate the individual's attributes (for example, the children's book bought in December was for a niece who was visiting, so further similar offers would be pointless).

SAS learned from the limitations of these self-learning algorithms, and developed the SAS Rapid Predictive Modeler. The SAS Rapid Predictive Modeler provides an environment for multiple models to compete within a simulation. The modeler automatically selects the best model, or adjusts the existing model to provide the highest lift. The SAS Rapid Predictive Modeler enables the next generation adaptive modeling environment—it blends low maintenance, while it facilitates control over the model, and achieves high lift. Real-time analytics can use the SAS Rapid Predictive Modeler to provide an automated high-performance modeling environment with continuous uptime for mission-critical environments.

STAYING RELEVANT

As customers interact with organizations across physical, digital, and social channels, it is critical to not only collect behavioral data, but to accurately identify individuals across channels, and to generate a holistic view of the customer and his or her circle of influence. Leading organizations using analytics are able to correlate individuals who research a product on the Web, and then make an in-store purchase. These same leading organizations analyze social networks to leverage the power of social sentiment. In other words, they manage key influencers that enhance and extend the reach of a product during product launches or promotional campaigns. In the near future, these leading organizations will leverage location-aware mobile applications, and pull customers into the store, based on the customer's proximity to the store location. These leading organizations will harness the power of social networks to pull in all of the customer's nearby friends as well!

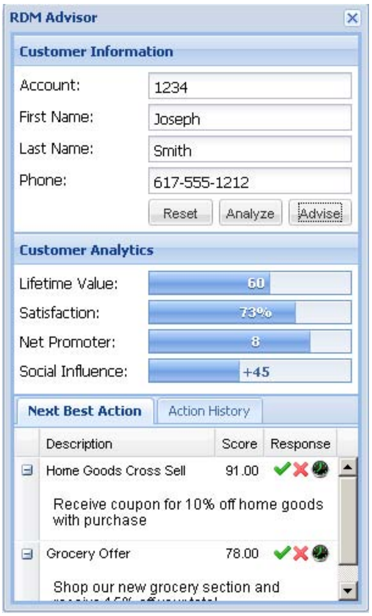

Figure 2~~3~~: Prototype UI for leveraging Real-Time Analytics and Next Best Action

CONCLUSION

Effective Next Best Action can increase both customer value and customer satisfaction by presenting relevant treatments during each customer interaction. It is most effective when real-time analytics are leveraged to provide up-to-date customer insight. Operationalizing analytics in rapidly changing environments has become feasible with SAS Real-Time Decision Manager and SAS Model Manager, by streamlining the procedural and technical challenges that have plagued IT and Business organizations

CONTACT INFORMATION

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

Chris Toth
SAS Institute Inc.
Boston, Massachusetts
E-mail: chris.toth@sas.com