Using Behavior Maps to Understand Customers

Next-Generation Technology in SAS® Interaction Management

The Power to Know®

A SAS White Paper

Table of Contents

Executive Summary

Technology can improve the way businesses communicate with their customers. Individual behavior can be profiled, monitored and used to trigger customer communications. With behavior-triggered communication, companies can move from product- and offer-centric monologues toward a two-way, customer-centric conversation. Businesses' ability to make this conversation engaging rests on their capacity to interpret the customer's underlying interests, needs and aspirations from observable actions.

This paper describes how to develop a better understanding of your customers based on their actions, reactions and inaction. We show how this understanding can dynamically change and how it can be used to drive real-time customer interactions. Our insights come from the unique description of the customer created from **Behavior Maps, a concept that is unique to SAS® Interaction Management.** These maps show marketers where customers have been, where they are, and where they are likely to go. Past behavior adds context and significance to current behavior, allowing marketers to anticipate customers' needs and problems and react to behavioral changes as they unfold.

A Conversation with the Customer

Every relationship between a business and a customer is best viewed as an ongoing conversation. Without some form of two-way communication, a business has no way to understand customer needs and problems, communicate solutions or set performance expectations. The better the customer conversation, the better the opportunity to maximize customer lifetime value.

For businesses with high profit contributions per customer (e.g., aircraft engine manufacturers/distributors), the customer conversation is no metaphor. It's the way those businesses acquire and manage customers — via personal sales calls, contract negotiations and post-sales relationship management. The conversation is a verbal communication between people, who deploy basic conversational skills (i.e., listening, asking questions, presenting ideas, interpreting body language and maintaining awareness of context) as a matter of course.

For businesses with large client bases and lower profit contributions per customer (e.g., telecommunications companies or retail banks), a face-to-face conversation is not feasible. The economics and scale of these businesses do not support establishment and maintenance of personal relationships. But the conversation does occur. Customers transact. They interact with businesses via automated touch points. Businesses make offers via direct marketing campaigns. Customers call to change or cancel relationships. These events constitute a potentially broad and deep conversation, even without direct verbal communication.

Customers are aware of their end of the conversation, because they experience a series of communications and understand in the context of the whole relationship. They know when a relationship is going well or headed for trouble. Do businesses know? Are they keeping up their end of the conversation?

The answer, in many cases, is no. Often when customers become dissatisfied with a relationship and begin to wind down their activities in favor of another provider, the business first becomes aware of the situation when the customer calls to terminate the relationship. Customers might receive (and businesses blindly deliver) the same direct marketing promotion several times in a row, regardless of whether or not the customer had previously responded. Customers are forced to continually re-introduce their needs and problems to a series of customer care representatives with limited capabilities and training.

Some part of a business is aware of every one of these non-verbal communications; they exist in transaction records, touch point logs and the like. That awareness is not unified and synthesized, however, in order to facilitate the conversation. Each new interaction typically occurs as an independent event, absent the context of the conversation preceding it.

The Importance of Context

Every layer of context adds knowledge. Knowing an individual's general preferences and interests is good. Knowing who a person is, what he does and what is going on in his life is better. Having this knowledge in light of what a person has done historically is best. Context allows us to recognize the relevance of specific behaviors and to anticipate future behavior.

Consider how a brokerage service might distinguish its customers. If a customer executes one trade in a particular month, the company might classify that customer as a low-level user. Based on this designation, the company might view the account as marginally profitable and choose not to spend money to protect it.

In contrast, if the customer had consistently traded ten or more times a month, then just one trade in a month sends a very different message. Rather than viewing the customer as a low-volume trader, the company sees a heavy trader whose behavior has changed. One possibility is that he has moved his trading elsewhere. This suggests the remaining activity and dollars might cease or be withdrawn, thereby placing the entire account at risk. By understanding the customer's current activity in the context of past behavior, marketers create an opportunity to respond in a relevant and timely fashion.

Creating Context

A participant in a conversation creates context from the words being spoken, body language of other participants and environment in which the conversation occurs. He is able to distill from this wealth of input data a sense of where the conversation is headed as well as develop a related understanding of actions (or non-actions) appropriate to achieve his own objectives (i.e., propound an idea, elicit some information, etc.).

While it is relatively easy for an individual to create context for a conversation, it is difficult for a business. Businesses must clear a number of hurdles to create context effectively:

- Input data must be collected from a variety of disparate sources.
- Data are too voluminous and varied to be useful in their raw form, so they must be reduced to a unified summary that is concise, organized according to actionable business concepts and readily available.
- Summaries must be created at the individual level, enabling the business to carry on thousands or millions of distinct customer conversations.
- Summaries must be continually updated with the latest transaction and interaction data, to ensure that the business is not a step behind in the conversation.

In recent years, businesses have built infrastructures to overcome these obstacles. Data warehouses and associated analytic tools have enabled businesses to collect large volumes of disparate customer data, create unified summaries of that data at the individual customer level (called profiles) and build mathematical models to predict customer behavior.

Are these businesses any better at holding up their end of the customer conversation? The answer is no. Data warehouses have been primarily used to drive targeted outbound marketing campaigns, focusing on optimization of offers rather than customer relationships. The results have been positive, from a marketing perspective — direct marketing campaigns have shown significant increases in response rates. But in the context of customer conversations, the results have been disappointing — an explosion of *non sequiturs*, solicitations with less than five percent probability of relevance for each individual customer, based on typical response rates to direct marketing campaigns. While data warehousing, analytics and optimization clearly are key enabling technologies for planning and strategic marketing, an outbound, offer-centric marketing strategy can be greatly improved by taking into account individual customer conversations when determining what offer should be made to whom, and when.

The challenge of creating context is not only a matter of focus but a matter of technology. Profiles in a data warehouse tend to be organized around input data (i.e., the database schema) and sophisticated mathematical constructs (i.e., predictive models), which are extremely valuable but not immediately comprehensible to the front-office personnel and automated applications responsible for carrying on the customer conversation. Also, profiles are typically updated periodically, not continuously. Due to the complexity and computational resource intensity of data warehouse loading procedures, most data warehouses are updated once a month. Consequently, the context is always out of date (except for that one time a month!), and the business is always a step behind in the conversation.

The Behavior Map

The Behavior Map and underlying technology was created to address these challenges by organizing data according to actionable business concepts and continuously updating the context for the customer conversation. One way to conceptualize a Behavior Map is as a moving picture, rather than the snapshot you would get from a static customer profile.

The Behavior Map is a representation of the customer that balances rich description with the ability to act. It is created using a set of non-overlapping states related to each other by behaviors. Past behavior determines what state a customer is in. Specific actions, or lack of actions, move that customer from the current state to a new one. The state provides the context in which all future behaviors can be interpreted.

A simple Behavior Map can be visualized as a directed graph. In Figure 1, we present an example capturing the trading behavior exhibited by a brokerage account. In this figure there are four states representing four levels of trading activity: dormant (0 trades), low (1-2 trades), moderate (3-5 trades) and high (>5 trades) (states are defined by global thresholds in this example for the sake of simplicity — in practice they are more commonly defined in relation to the individual's "normal" activity).

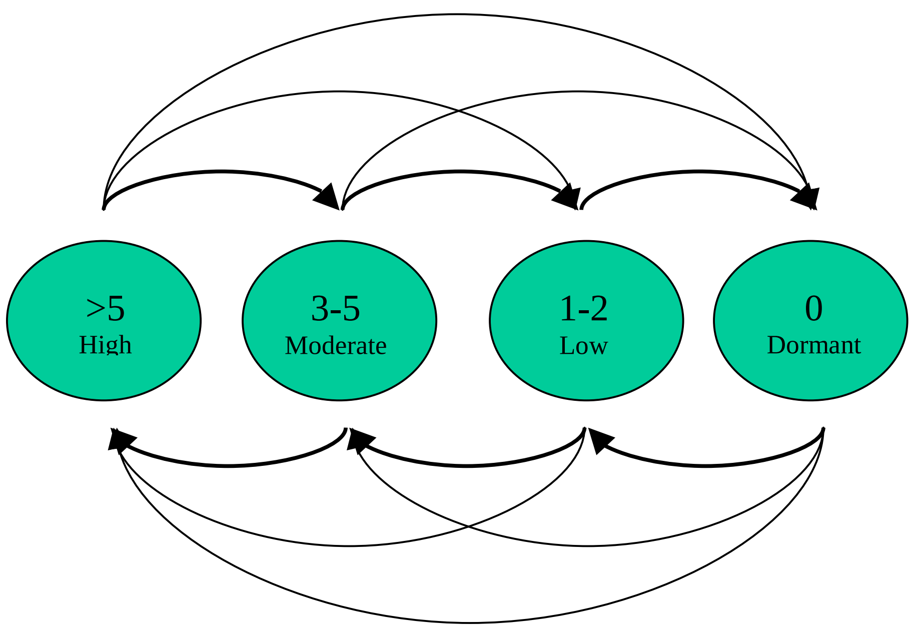

Figure 1: A Behavior Map describing activity volume.

A customer's location on a map, or his current state, reflects recent activity as it relates to prior behavior. In Figure 1, for example, the number of trades made in the past month provides behavioral context for how recent behavior positions the customer on a map. A customer who had made more than five transactions a month and suddenly changes to zero moves along a particular path, from the state at the far left of Figure 1 to the state at the far right.

Actions can be associated with a particular transition or sequence of transitions. As noted above, a company would probably want to treat customers differently who, in a single transition, move from high to dormant activity than those who move from high to moderate activity.

Behavior Maps offer a view of customer behavior that encompasses inactivity as well as activity. A state can be defined by what is not done. If an individual applies for a credit card, but does not use the card in the first 30 days, that inactivity is sending a message. There may be some feature of an existing card that the cardholder was not willing to give up or there may be a problem with the card activation process. A Behavior Map can be designed to capture and react to this inactivity.

Additional context is captured as a behavior path — the sequence of states in a Behavior Map that an individual customer has visited over time. These paths integrate historical behavior with the customer's current state. Behavior Paths tell us not only where customers have been in their conversations, but also how long they resided in each state along the way. One customer's trading volume might increase to three per month over a few months, while a second customer might take more than a year to reach the same volume. The first customer is exhibiting significantly faster growth and may hold greater future potential. By capturing the length of time an individual resides in each state, it is also possible to determine the velocity at which the customer's behavior changes. This information can help identify trends, volatile behavior and cyclic behavior.

Paths within a Behavior Map can be determined by the actions of the business as well as the customer. A series of service delivery failures by a service provider can increase the likelihood of account closure, for example. This change can be represented by transitioning the account into a state that is characterized by higher risk of attrition. A variety of actions can be associated with this state to win back the customer's confidence.

Acting on Behavior Maps

A Behavior Map provides a natural integration of description and action. It is designed to (1) differentiate "normal" from "abnormal" behavior and (2) classify "abnormal" behaviors as potential opportunities or threats, where "normal" is established and continuously monitored at the individual level. Business rules can be attached to each classification, directing interventions to take advantage of opportunities and avert threats. Equally important, the context for interventions can be transmitted to the field as concise descriptions of the "abnormal" behavior and why it may constitute an opportunity or threat.

This approach is the basis for credit card fraud detection, where the goal is to detect when abnormal or suspicious activity occurs. To detect credit card fraud, you start by building a profile of normal behavior for each customer. When the customer's behavior deviates from this baseline state, the customer transitions to a suspicious state. This deviation can be caused by changes in transaction velocity, such as an unusually large purchase or purchases at non-conforming merchants. Because only two states are used, the simple detection of change may be sufficient to initiate a review by a fraud investigator. Likewise, by defining the state change on a relative basis instead of an absolute volume of transactions, false positives are avoided, and customers who normally have high transaction volumes are not erroneously classified in the suspicious state. The practical results of relative state changes are clear—investigators don't waste time on false alarms and can instead focus on truly relevant alerts; and innocent customers are less likely to be investigated for fraud. Fraud detection costs (both in terms of actual overhead and fraud-related losses) are reduced through this method, and customers are better protected against fraud.

In some situations, additional information from a Behavior Map can help to refine context for intervention, improving the chances for relevant participation in the conversation. For example, knowing someone is price sensitive can suggest what message is most appropriate to stimulate a sale. Knowing the preferred channel can indicate how to send that message. In other situations, such as in the brokerage example above, marketers might choose to consider environmental factors, such as overall market activity, before reacting to an individual's drop in activity. Likewise, they might wish to incorporate the results of more complex analyses, such as a "propensity to buy" JScore created by SAS Enterprise Miner, in determining the next step in the customer conversation.

Combining Simple Behavior Maps

Another possibility is to build individual Behavior Maps for homogenous sets of transactions to create multiple views of the same customer. For example, to describe a customer's relationship with her bank, we could build one Behavior Map based upon her deposit activity and others on withdrawal activity, credit card transaction volume or fees assessed.

In such cases, an individual will reside simultaneously in one state of each map. For example, the individual might be in a state showing increased transaction behavior on one map and at the same time be in a state on a different map showing decreased deposit activity.

One Behavior Map may be a strong predictor of a specific event, but there could be other indicators as well. So, for example, one sign that a brokerage account is about to attrite may be a drop in account activity. Another sign might be large out-of-pattern withdrawals. Although these states might reside in two distinct Behavior Maps, a customer's presence in either state could be identified. The logical combination of the members of one state with the members of another state, either in the same or a different map, is how more specific or broader populations are defined.

Tactical and Strategic Behavior Maps

The concept of customer Behavior Maps can also be extended to support a long-term strategic vision of the customer relationship. Behavior Maps are expressive enough to support a complete view of the customer from cradle to grave. They can provide a road map for maximizing the lifetime value of the customer from the initial solicitation, through enrollment and activation to cross selling and retention. An abbreviated example of one such map is shown in Figure 2.

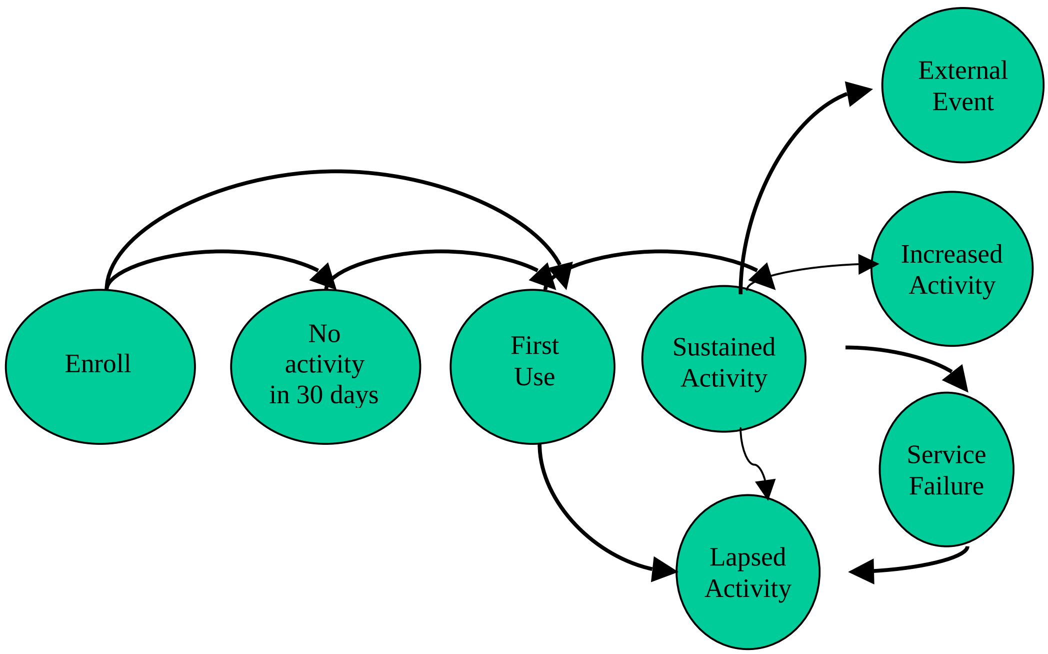

Figure 2: A representative Behavior Map that could be used to strategically plan the customer experience.

Data that Support a Deeper Understanding of the Customer

Up to now, we have focused on the unique advantages of organizing the data in a Behavior Map. In this section, we will discuss some of the data sources used in customer analysis today and how these might change to take advantage of the Behavior Map.

In today's environment, direct marketing tends to use descriptions of the customer that do not vary with time. As a result, they rely on broad demographic segments. In some cases, geographic averages are substituted when individual data does not exist. The use of such averages is based on broad generalizations and stereotypes. The first priority of moving to a customer-centric conversation is to move beyond broad profiles to more dynamic individual profiles.

One method of capturing an individual's evolving interests and needs is to ask. This is the approach taken by warranty cards and Web sites that require extensive registration surveys. Users must answer dozens of questions gathering detailed demographic and financial data as well as interests in a variety of fields. Even if you succeed in getting a user to fill out the profile honestly, the data still only represent a snapshot in time. New customer representations must be more dynamic if they are going to drive relevant messages.

An alternative approach is to develop a picture based on observing and interpreting behavior. This transaction analysis requires minimal participation by the user and provides a rich description of the "real" conversation occurring between the business and the customer. This is the approach taken by rewards clubs and frequent shopper clubs, where customers are given discounts in exchange for recording their detailed purchases. Clickstream data are similar to financial transaction in that they unobtrusively provide insight into behavior. These data may identify what an individual looks at, where they choose to linger and to where they return. Through such data, companies can understand the context required to carry on their end of the customer conversation.

Enhancing a Behavior Map with Other Descriptors

In some cases observed changes in behavior are particularly relevant in light of other events. Consider a credit card example, in which the Behavior Map is oriented towards predicting customer attrition. In this map, a drop in transaction volume may move a credit card holder from a high-volume transactor state to just a moderate-volume state. Since there are other cardholders who consistently demonstrate moderate transaction volume, the cause for concern is the change in velocity rather than the velocity itself. Conversely, global external drivers can be readily deseasonalized by redefining the “high transactor” state relative to aggregate behavior. In the scenario above, if the drop in transaction volume occurred immediately after the holiday season, the individual cardholder whose transaction volume decreased may stay in the high transactor state, provided her behavior was within the seasonally adjusted definition of that state.

Another reason to combine external events with Behavior Maps is to anticipate behavioral changes. Significant life events such as graduation, marriage, birth of a child, death of a relative or personal illness can dramatically change an individual’s needs and behaviors. Knowing about these events before they change observable behavior allows you to proactively address the customer’s needs. Attributes derived from behavioral data or identified from subscription lists, memberships lists and self-reported registrations can help refine the message. The attribute values can be used to define the specific content, features and copy of the offer.

Behavior Maps can also provide businesses with a "customer's eye view" by looking at the performance of their own products and services relative to the markets in which they compete. With this understanding, organizations can track even subtle changes in customer sentiment by understanding how every individual client perceives their products. For example, a financial services company could anticipate their customer's satisfaction by knowing how a customer's investment portfolio stacks up against an industry average. Likewise, a wireless telecom provider could better project customer churn for their individual customers by comparing their own local pricing plans to those of their competitors.

Conclusion

In this paper we have presented a new way of describing the customer that will allow businesses to optimize the value of customer relationships by engaging the customer in a conversation. We have shown how Behavior Maps provide unique capabilities for detecting and reacting to changes in customer behavior.

Customers want to be listened to and treated with respect. For those businesses that cannot meet the challenge, there are competitors who will. And in the end, another offer may be only a click away.

Appendix A — Behavior Maps in Relation to Traditional Modeling

Behavior Maps and Other Models

Behavior Maps complement the SAS predictive models that are used ubiquitously in direct marketing. To appreciate how they work together, we will review how predictive models are currently used and where they are challenged.

Predictive models are commonly built using regression, decision trees or neural networks. These models compute the expectation of a specified behavior based upon a set of account attributes. For example, by looking at the attributes of one's credit bureau, it is possible to predict the likelihood of that person going bankrupt within the next 24 months. In most cases, predictive models are used to rank order individuals for the purpose of managing resources. By rank ordering the individuals, the number of direct mail contacts can be limited to live within the constraints imposed by a budget, operational capacity or overall risk exposure.

There is inherent ambiguity associated with the scores that are generated by predictive models. The specific probability of an event associated with a score is dependent upon the population used to build the model. If other criteria or scores are used to segment the population, then the probabilities for a given score change. Consider a model that is built to predict the likelihood of an individual going bankrupt in the next 12 months. If a response model is used to reduce the size of a prospective pool of credit candidates, then the underlying population introduces a bias. Individuals with the same bankruptcy score are no longer a random sample of such scores drawn from the full population. Consequently, the probability of bankruptcy associated with that risk score has changed. In fact, because those likely to respond to a new offer for credit tend to be riskier, the odds of default for the same score are higher than before.

Addressing time with predictive models poses a challenge. A credit-bureau-based acquisition risk model attempts to identify individuals that will default in the next 24 months. The model is asked to use the same set of attributes to identify those individuals who will default tomorrow as well as those who default in 24 months. As a result, the model looks at the average impact of specific events over the duration of the observation window. This blurring of time makes it difficult to exploit the occurrence of specific events that tend to precede financial disaster.

Capturing the impact of context is also difficult with predictive models. An individual's credit bureau risk score can provide a general description of a customer's likelihood of defaulting. However there are many factors that contribute to that score. A low score could be a result of being young with no track record or being older and historically delinquent. To a credit lender, these are two very different messages, and the only way these differences can be illuminated is by providing context.

Using Behavior Maps to Complement Predictive Models

There are two characteristics of the Behavior Map that help traditional predictive models handle time and context. Behavior Maps segment populations into distinct groups based upon historical behavior (actions *and* non-actions) in time. This segmentation can be used to build more accurate models and also improve the accuracy with which old models are interpreted. Secondly, Behavior Maps can accomplish this segmentation with little or no data. Therefore, when a given model cannot be used because all of the required input data is not available, a Behavior Map can provide an alternative means of identifying individuals of interest.

Segmentation can help build better models by dividing up the population into groups with more homogenous behavior. Consider a model being built to predict the risk of default on a given credit card. For those cardholders who historically never miss a payment and always pay on time, a sequence of late or missed payments is an immediate sign of trouble. In contrast, there are those who are sloppy payers, who may travel and miss a payment or are just inattentive of their finances. For these cardholders, missing a payment does not communicate the same message. If both of these populations were in the same development set, then the model would not be able to use that information. Since Behavior Maps capture individual patterns, they are able to discriminate between the populations.

For similar reasons, segmentation can also be used to increase the level of confidence in a particular model score. By calibrating the performance of the model score within each state, users can identify when the score is most predictive and when the model breaks down. This information allows the models to be used more selectively and helps reduce the margin of error around the scores when they are used.

By combining Behavior Maps, predictive models and other SAS analytic technologies, SAS Interaction Management clearly provides a robust architecture for delivering valuable customer intelligence.

World Headquarters
and SAS Americas
SAS Campus Drive
Cary, NC 27513 USA
Tel: (919) 677 8000
Fax: (919) 677 4444
U.S. & Canada sales:
(800) 727 0025

SAS International
PO Box 10 53 40
Neuenheimer Landstr. 28-30
D-69043 Heidelberg, Germany
Tel: (49) 6221 4160
Fax: (49) 6221 474850

www.sas.com

100693US_221861.0103