

IBM Predictive Customer Intelligence
Version 1.0.0

Solution Guide



Note

Before using this information and the product it supports, read the information in "Notices" on page 87.

Product Information

This document applies to IBM Predictive Customer Intelligence Version 1.0.0 and may also apply to subsequent releases.

Licensed Materials - Property of IBM

© **Copyright IBM Corporation 2014.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Introduction	vii
Chapter 1. IBM Predictive Customer Intelligence	1
Create a Predictive Customer Intelligence solution.	2
Chapter 2. Case studies	3
Telecommunications case study	3
Use Predictive Customer Intelligence to retain a dissatisfied customer	4
Define the offers that customers of a telecommunication company can receive	7
Determine the best offers for telecommunications customers by creating business rules	8
Build predictive models for a telecommunications company	10
Monitor the success of offers	10
Retail case study.	11
Predictive retail promotions process	11
Use Predictive Customer Intelligence to increase revenue through targeted recommendations	12
Define the offers to be presented to retail customers.	13
Determine the best offers for retail customers by creating business rules	14
Build predictive models for retail promotions	15
Use Predictive Customer Intelligence to build cross-channel loyalty: in-store promotions	16
Insurance case study	17
Use Predictive Customer Intelligence to drive cross-sell opportunities	17
Define the offers to be presented to insurance customers	18
Determine the best offers for insurance customers by creating business rules	20
Build predictive models for an insurance company	20
Energy and Utilities case study.	21
Banking case study.	22
Use Predictive Customer Intelligence to target promotional offers to the right customer	22
Define the offers that banking customers can receive	23
Build predictive models for a banking customer	24
Chapter 3. Predictive models	25
Data sources	25
Predictive models in the Telecommunications sample	26
Predicting churn.	27
Customer satisfaction	28
Assigning offers	29
Targeting offers to customers with the response propensity model	30
Telecommunications models in Analytical Decision Management	30
Predictive models in the Retail sample	31
Prepare data for a retail solution	31
Defining customer segments.	33
Market Basket Analysis	34
Determine Customer Affinity	35
Response log analysis	36
Inventory-based suggestion	36
Deployment of models in the Retail sample	37
Use Retail case study models in IBM Analytical Decision Management	39
Predictive models in the Insurance sample	40
Data used in the Insurance sample	41
Define customer segments in the Insurance sample	42
Predict churn in the Insurance case study	42
Understand Customer Lifetime Value (CLTV)	43
Predict customer response to a campaign	44
Predict churn for auto policy holders.	44
Group customers based on their current life stage	45

Identify the life policies bought by life stage segments	45
Recommend the correct insurance policy	45
Transform and aggregate customer data	45
Extract life stage event information from social media posts	46
Extract sentiment scores from customer complaints	46
Insurance data model	46
Predictive models in the Energy and Utilities sample	48
Predict credit rating	48
Predict the need for demand response program assistance	49
Recommend the right rate plan	49
Understand customer sentiment	49
Predictive models in the Banking sample	49
Determine category affinity in banking	50
Predict churn in the banking industry	50
Predict whether customers will default on credit card debt	50
Define customer segments in the banking industry	51
Sequence analysis	51
Training predictive models	51
Scoring a model	51
Creating business rules	52
Deployment	52
Chapter 4. Integration with the IBM Enterprise Marketing Management suite.	53
IBM Interact connectors	53
The External Learning connector	53
The External Callout connector	54
Chapter 5. Integration with Next Best Action Optimizer	55
Open NBAOPT Studio from IBM SPSS Modeler	56
Generating streams from PMML files	56
Example SPSS stream	57
Configuring the IBM SPSS stream	58
Chapter 6. Reports and portal pages	59
Sample reports and portal pages	59
Viewing the sample reports	61
Modifying the Predictive Customer Intelligence data model	61
Predictive Customer Intelligence Usage Report	62
Running the Predictive Customer Intelligence Usage Report	63
The Framework Manager model for the Usage Report	63
Appendix A. IBM Predictive Customer Intelligence samples	67
Telecommunications sample	67
Retail sample	68
Insurance sample	69
Banking sample	70
Energy and Utilities sample	71
Appendix B. Configuration of the Usage Report	73
Predictive Customer Intelligence database tables	73
Configuring logging in IBM Enterprise Marketing Management	76
Populate the Predictive Customer Intelligence database from Enterprise Marketing Management	77
Configure logging in IBM SPSS Collaboration and Deployment	78
Populate the Predictive Customer Intelligence database from IBM SPSS	80
Appendix C. Troubleshooting a problem	83
Troubleshooting resources	83

Notices	87
Index	91

Introduction

IBM® Predictive Customer Intelligence gives you the information and insight that you need to provide proactive service to your customers. The information can help you to develop a consistent customer contact strategy and improve your relationship with your customers.

Audience

This guide is intended to provide users with an understanding of how the IBM Predictive Customer Intelligence solution works. It is designed to help people who are planning to implement IBM Predictive Customer Intelligence know what tasks are involved.

Finding information

To find product documentation on the web, including all translated documentation, access IBM Knowledge Center (<http://www.ibm.com/support/knowledgecenter>).

You can also access PDF versions of the documentation from the Predictive Customer Intelligence web page (www.ibm.com/support/docview.wss?uid=swg27041723).

Accessibility features

Accessibility features help users who have a physical disability, such as restricted mobility or limited vision, to use information technology products. Some of the components included in the IBM Predictive Customer Intelligence have accessibility features.

IBM Predictive Customer HTML documentation has accessibility features. PDF documents are supplemental and, as such, include no added accessibility features.

Forward-looking statements

This documentation describes the current functionality of the product. References to items that are not currently available may be included. No implication of any future availability should be inferred. Any such references are not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of features or functionality remain at the sole discretion of IBM.

Samples disclaimer

Sample files may contain fictional data manually or machine generated, factual data compiled from academic or public sources, or data used with permission of the copyright holder, for use as sample data to develop sample applications. Product names referenced may be the trademarks of their respective owners. Unauthorized duplication is prohibited.

Chapter 1. IBM Predictive Customer Intelligence

IBM Predictive Customer Intelligence brings together in a single solution the ability to do the following tasks:

- Determine the best offer for a customer.
- Retain customers that are likely to churn.
- Segment your customers, for example, by family status and salary.
- Identify the most appropriate channel to deliver an offer, for example, by email, telephone call, or application.

This solution ensures that all interactions with customers are coordinated and optimized. Predictive customer intelligence gives you the ability to sift quickly through millions of subscribers and know who to contact, when, and with what action.

The following steps define the process:

1. Understand the customer. Predictive modeling helps you to understand what market segments each customer falls into, what products they are interested in, and what offers they are most likely to respond to.
2. Define possible actions and the rules and models that determine which customers are eligible for which offers.
3. After the best action is identified, deliver the recommendation to the customer.

Optionally, you can monitor the effectiveness of your solution by using the Usage Report that is provided. The Usage Report displays the number of offers that are presented to customers and can be configured to show the number of offers that are accepted and rejected. For more information, see “Predictive Customer Intelligence Usage Report” on page 62.

Integration with the IBM Enterprise Marketing Management (EMM) suite

IBM Predictive Customer Intelligence integrates with the following solutions:

- IBM Campaign, a web-based solution that enables users to design, run, and analyze direct marketing campaigns.
- IBM Interact, which provides personalized offers and customer profile information in real time.
- IBM Marketing Platform, which provides security, configuration, and dashboard features for IBM EMM products.

IBM Predictive Customer Intelligence provides two connectors between IBM Interact and IBM SPSS® Collaboration and Deployment Scoring Service:

- The External Callout connector calls an IBM SPSS model at run time, and is contained within the expression of an advanced rule for a marketer score, overriding the score that is supplied by the EMM campaign.
- The External Learning connector extends IBM Interact's native learning module to monitor visitor actions and propose optimal offers. It prioritizes IBM Campaign offers based on an IBM SPSS model's prediction of their final score. The connector passes specific configurable parameters as input to the IBM SPSS Scoring Service.

Integration with the IBM Next Best Action Optimizer

IBM Next Best Action Optimizer (NBAOPT) segregates customers by lifetime value, and then provides recommended actions to retain customers, which are based on their lifetime value segment.

You can integrate NBAOPT with the IBM Predictive Customer Intelligence solution. Install the NBAOPT connector to add an NBAOPT item to the IBM SPSS Modeler **Tools** menu. You can start the NBAOPT Studio from this menu, and you can generate a stream from PMML files.

Create a Predictive Customer Intelligence solution

Sample files are supplied to help you to design an IBM Predictive Customer Intelligence solution for the needs of your business.

Samples for the following industries are included:

- Telecommunications
- Retail
- Insurance
- Energy and Utilities
- Banking

Case studies that describe common scenarios for each industry where Predictive Customer Intelligence can be used are provided. See Chapter 2, “Case studies,” on page 3.

For information about installing the samples, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

Chapter 2. Case studies

Case studies illustrate common scenarios where IBM Predictive Customer Intelligence can be used to provide prioritized offers to your customers.

Telecommunications case study

Telecommunications providers are leaders in harnessing big data, for example: billing data, demographic data such as gender, age, and employment status, transaction history, call detail records, and call center records.

The key to success is bringing different data sources together, such as structured data, unstructured data, internal data, and external data, to create a profile of each customer. By constantly analyzing all of these information types, IBM Predictive Customer Intelligence provides you with insights that you can use to build a service that is fine-tuned to a customer's specific needs.

This case study demonstrates how IBM Predictive Customer Intelligence can be used in the telecommunications industry to retain a dissatisfied customer:

Retain a dissatisfied customer

A call center agent uses the data in the call center application to do the following activities:

1. Recognize that the customer is dissatisfied through analysis of network traffic and call detail records.
2. See that the customer is high value.
3. See what actions and offers the customer is eligible for.
4. Proactively submit an action for a service level request to fix issues, and contact the customer with an offer.

Determine which offers customers receive based on profile and real-time data

Using IBM Campaign, the marketing manager determines which offers customers can receive based on their customer profile data and their real-time interaction data. For example, a new product inquiry, a complaint, or a data plan inquiry.

Determine the best offers for customers by creating business rules

Using IBM Analytical Decision Management, the business analyst creates business rules to determine which actions are valid for a customer. For example, you might create a rule that targets retention actions to customers with high network influence, high churn risk, and no open cases.

Predict churn, customer satisfaction, and propensity to respond to offers by creating predictive models

Using IBM SPSS Modeler, a data modeler creates predictive models to predict the following factors:

- A customer's propensity to churn.
- Customer satisfaction.
- Propensity to respond to offers.

The results from these models are used in Analytical Decision Management.

Monitor the success of offers

Marketing managers and business analysts can monitor the success of offers with the Usage Report. The Usage Report provides a snapshot of how many offers were presented to customers by channel and by month.

Use Predictive Customer Intelligence to retain a dissatisfied customer

Bill is the customer of a telecommunications firm and is experiencing excessive dropped calls and handset issues. This case study describes how IBM Predictive Customer Intelligence can be used to help to retain Bill, who is a high value customer.

This example demonstrates how IBM Predictive Customer Intelligence coordinates the telecommunications company's interactions with its customers, ensuring that customers are contacted at the right time, through the right channel, and with the best possible actions. It uses the following architecture: The call center application connects directly to IBM Enterprise Marketing Management (EMM) by using the IBM EMM API.

Customer profile

Bill is a professional architect and a small business owner. He is married but with no children. He has a single mobile account with the telecommunications provider that he uses both for business and personal use.

Bill is an extensive social and mobile application user. He does extensive web browsing and makes many telephone calls.

Bill is a high value customer, one the company wants to retain. However, Bill is dissatisfied. Analysis of network traffic and call detail records show that Bill recently had a number of dropped calls, he called a competitor's call center, and he has two hardware issues with his phone.

Carolyn, a call center agent initiates an outbound call to Bill. She opens her Customer Relationship Management (CRM) software and looks at Bill's profile. From a dashboard interface, she sees three alerts that require her attention:

- Bill is eligible for support.
- He is eligible for offers.
- Bill is at great risk of leaving the company.




Figure 1. Dashboard notifications that shows alerts

Carolyn can also see some metrics for Bill. The data is provided by IBM SPSS predictive models and displayed in Cognos® Business Intelligence reports:

- His churn score is High.
- Customer Satisfaction is Low.
- Customer Lifetime Value is High.
- Social Influence is High.



Figure 2. Dashboard profile that shows metrics

These metrics indicate that Bill is a valuable and influential customer who is at risk of leaving.

Carolyn can see that Bill has a history of issues with his phone. She can see through real-time sentiment analysis of Twitter feeds that Bill is experiencing problems with dropped calls, which he has not yet reported to the service center. Real-time detection of dropped calls is causing the dropped call indicator to reach a critical threshold, which the agent sees identified.

Bill is not happy with his current phone. The company needs to intervene to retain Bill as a customer.

Actions

Carolyn looks up at the Notification area at the top of her CRM screen to see the available Actions for Bill.

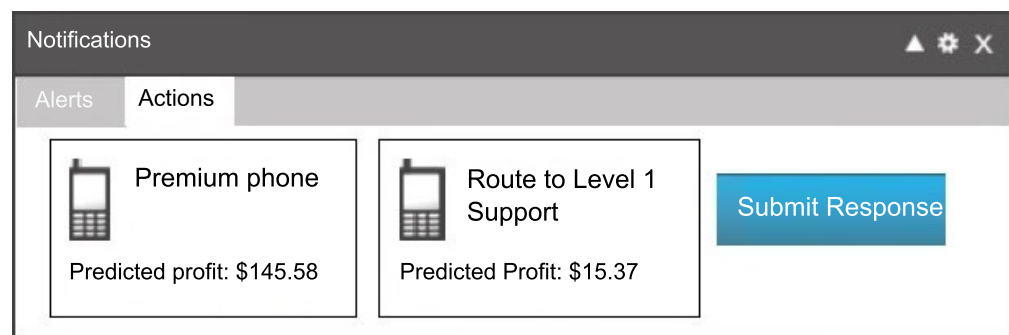


Figure 3. Notifications area that shows actions

The second action is a proactive service request that can be submitted by Carolyn on behalf of Bill. This action addresses the dropped call issue that Bill experienced. Carolyn selects the support action and clicks **Submit Response**.

Carolyn tells Bill that she is submitting a Level 1 service request on his behalf to fix the dropped call issues that he experienced. She tells him that the issues will be fixed immediately so that he can continue to have uninterrupted service. Bill is impressed that his service provider detected the issue without him calling in to report it, and took proactive steps to have the dropped call issues fixed.

Due to Bill's recent hardware issues, Carolyn also offers Bill a free upgrade to a premium phone. Bill readily accepts the offer of a new phone.

After Carolyn submits his response, Bill's scores are updated. His Churn score now shows a reduced risk. It is now at 55%, which is down from the previous 75%. Both his Satisfaction and Customer Lifetime Value scores increased.



Figure 4. Dashboard profile that shows updated metrics

Optimization of offers

You can further optimize offers. Following on from the previous example, connect the call center application to IBM Analytical Decision Manager, and then connect Analytical Decision Manager to your existing IBM EMM campaigns by using the External learning connector.

The External learning connector provides different offers for Bill, under the same circumstances. He is offered:

- Route to Level 1 support.
- Initiate Network Service Request.

These options are less expensive to the Telecommunications company, but are equivalent in reducing churn and increasing satisfaction.

Offer optimization workflow

The following steps demonstrate how IBM Predictive Customer Intelligence determines the offers to be presented to Bill by using the External Learning connector.

1. Open customer profile

Carolyn, the call center agent, types Bill's phone number in the Customer Relationship Management (CRM) system and clicks **Search**.

2. Obtain actions and alerts

The CRM system passes a set of parameters to the IBM Enterprise Marketing Management (EMM) Interact Web Service as a request using the Interact API. The Interact web service outputs the offers.

3. Personalize offers

The Interact web service passes a set of parameters including the offers to

the IBM SPSS Scoring Service using the External Learning connector. The External Learning connector returns a score for each offer to the call center agent system.

4. Reorder offers

The CRM system sorts the list of offers using the score. The CRM system then calls the IBM SPSS Scoring Service with the list of offers and other customer parameters using the Scoring Service API. An IBM Analytical Decision Management model is activated.

5. Get more offers

IBM Analytical Decision Management issues new offers to the CRM system through the IBM SPSS Scoring Service. The offers are displayed in the CRM system.

6. Run reports

The CRM launches the report that is generated by IBM Cognos Business Intelligence.

7. Submit customer response

Carolyn presents the offers to Bill, and clicks **Submit Response** when he accepts.

8. Re-run models

The call center agent system passes a set of parameters to the IBM SPSS Scoring Service as a request using the Scoring Service API. The IBM SPSS web service re-runs the applicable models and outputs the parameters of the model results. Bill's customer profile is updated in the database.

9. Obtain actions and alerts

The CRM system passes a set of parameters to the Interact web service as a request using the Interact API. The Interact web service outputs the offers. If there are no further actions, the process stops here. If there are any further actions and alerts, the process goes back to the Personalize offers step.

Define the offers that customers of a telecommunication company can receive

Monica is a marketing manager who is responsible for defining her company's marketing strategy. She determines which offers individual customers can receive based on their customer profile data and their real-time interaction data.

Customer profile data and real-time data

The customer profile data can include the following:

- Customer behavioral data, call, and text volume, products owned, contract details.
- Customer demographic information and contact preferences.
- Predictive model scores.
- Social data.
- Sentiment.
- Prior transactions and campaign responses.

Real-time data can be found in the context of the current call, such as a new product inquiry, a complaint, or a data plan question.

The campaign

Monica creates a campaign that is named Telco Call Center Campaign by using IBM Campaign, part of the IBM Enterprise Marketing Management (EMM) suite.

The Telco Call Center Campaign contains offers to be presented within the call center application. The details of the offer vary based on individual customer characteristics. The results of the offer, whether accepted or rejected, are recorded in the customer's interaction history, and that information can be used for follow-up marketing offers and campaigns.

Customers are segmented into categories in real time, based on the context of the call, real time scoring, and the campaign's business rules.

For example, you might have a flowchart with a decision node that assigns customers to five categories:

- HighValue
- ProductInquiry
- Complaint_LowValue
- Complaint_HighValue
- CloseAccount

In this example, customers are assigned to the HighValue segment if they have a customer lifetime value ratio greater than 0.7.

Each segment is assigned specific offers. Offers can be enabled or disabled by the marketer as the marketing strategy changes.

The HighValue customers qualify for the following potential offers:

- Premium phone.
- Premium data plans.
- Premium family plan.
- A second line at no cost.

Offer eligibility criteria for customers is based on calculations that use the real-time customer data and offer definition attributes.

After each offer is evaluated, there might be more offers eligible than there is time or space to present. For example, a customer is eligible for three offers, but you want to present only the top two offers. In this case, the offers would be presented based on the marketer's score values. The marketer's score can be based on a calculation of an expression by using customer and offer attributes, a real time learning algorithm, or a combination of these methods.

Determine the best offers for telecommunications customers by creating business rules

Alex is a business analyst. He uses IBM Analytical Decision Management to bring together his company's business rules, predictive models, and optimizations to determine the best possible offer for a customer.

Alex creates business rules to determine which actions are valid for any single customer. Some examples of what you might set up business rules for are shown in the following list:

- Retention
- Billing
- Service
- Support
- Marketing

The telecommunications company has a retention campaign. Alex targets retention actions to customers with high network influence, high churn risk, and no open cases.

Offers are allocated so that the first offer in the list that applies is presented to the customer.

For example, customers are offered a premium phone if the following is true:

- Offer1 equals premium phone. This value was assigned by IBM Enterprise Marketing Management.
- Current_offer does not equal premium phone. The current offer is what was on record in the customer's profile, generated from the last time an offer was made to the customer. You do not want to offer the same thing to a customer twice.
- The association model is a segmentation model that determines whether a customer is a promoter or a detractor, and whether the customer is amenable to offers.

If a customer qualifies for multiple offers, these offers can be prioritized.

You can set the maximum number of offers (**Max number of offers**) to be presented to each customer, for example, two.

You can apply an optimization equation to each action that is valid for a customer. The aim in the example is to pick the top two actions. When the equation is applied, the actions with the highest two scores are presented to the customer.

$RESPONSE \text{ PROPENSITY} * (REVENUE + (LOYALTY * (CHURN * FUTURE \text{ REVENUE})))$

The equation contains the following key attributes:

- Propensity to respond to the offer.
- Revenue the offer would bring in.
- How much impact the action would have on customer loyalty and churn.
- Predicted future revenue.

Not all data needs to come from rules that are created in IBM Analytical Decision Management. Some data can be defined elsewhere and imported into Analytical Decision Management. The key is to bring the intelligence together in one place and to orchestrate the final decision within a single application:

- Predictive models can be built in another tool by a data miner.
- Marketing campaigns and offers can be defined by a marketer.
- Data can be made available in the right form for modeling and deployment by an information architect.

For more information, see the IBM Analytical Decision Management Application User's Guide (http://www-01.ibm.com/support/knowledgecenter/SS6A3P_8.0.0/com.ibm.spss.dm.userguide.doc/configurableapps/user_overview_container.htm).

Build predictive models for a telecommunications company

David, the data scientist (or data engineer), builds models using IBM SPSS Modeler to make the following predictions for a customer:

- The propensity to churn.
- Customer satisfaction.
- The propensity to respond to offers.

The scores, segments, and results from these models are used as inputs within IBM Analytical Decision Management by Alex, the business analyst, as he creates business rules.

This is an example of a model that uses data from a corporate database to compute a churn score from 0-100% for each customer.




Figure 5. A model that computes churn scores for customers

This is an example of an Association model. An Association model determines whether a customer is a promoter or a detractor, and whether the customer is amenable to offers.

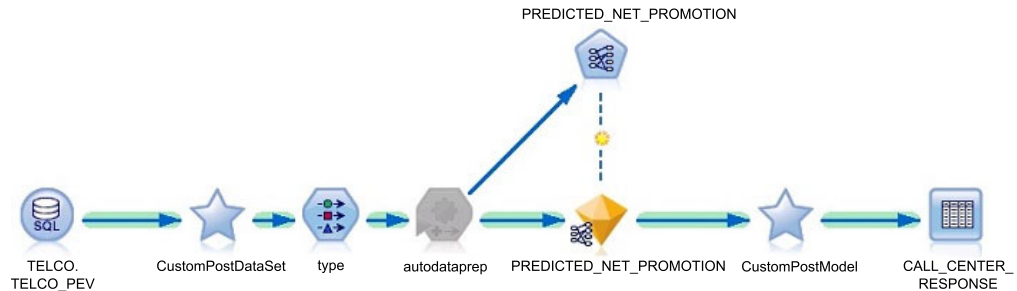


Figure 6. Example of an Association model

For more information, see “Predictive models in the Telecommunications sample” on page 26.

Monitor the success of offers

The telecommunications company wants to know how well their offers from IBM Predictive Customer Intelligence are performing.

They want the answers to the following questions:

- How many offers from IBM Predictive Customer Intelligence are being given to customers?
- How many offers are being accepted and how many are being rejected?
- How can this information be presented to executives?

The Usage Report provides a snapshot of how many offers were presented to customers by channel and month. This data can be filtered by date, response type, and channel, and can be used to determine and present the effectiveness of the offers. For more information, see “Predictive Customer Intelligence Usage Report” on page 62.

Retail case study

This case study demonstrates how IBM Predictive Customer Intelligence can be used in the retail industry to achieve the following results:

Understand the customer:

Profile customers based on browsing and purchase history.

Track propensity to respond to different types of offers.

Identify the behaviors that convert browsers to shoppers.

Find the right offer:

Increase profits with targeted offers informed by market basket analysis.

Combine the customer profile and history to identify the next likely purchase.

Prioritize offers based on expected profit.

Find the right channel:

Tailor offers to the point of interaction.

Build cross-channel loyalty with coupons informed by online activity for use in the physical store.

Predictive retail promotions process

Use predictive modeling to understand what market segments each customer falls into, what products they are interested in, and what offers they are most likely to respond to.

Then, you can identify the stage of the online purchase funnel that the customer is in:

- visitor
- product viewer
- buyer
- repeat buyer

You can determine how valuable the customer is in terms of revenue, and look at what the demographic profile says about the customer. For example, would they be interested in purchasing Product X, based on past purchase behavior? You can also establish the optimal discount that will entice a customer to purchase the product.

Using IBM SPSS Modeler, you can build models that predict how a customer is likely to behave.

Define possible actions

The next step is to define the rules and models that determine which customers are eligible for which offers. The goal is to determine how a customer can be

progressed through the online purchase funnel: that is, from a visitor to a product viewer, a product viewer to a buyer, and a buyer to a repeat buyer.

Using IBM Analytical Decision Management, you can combine predictive models with business rules to allocate offers in accordance with business goals.

Deliver best actions to the point of interaction

When the best action is identified, the recommendation is delivered to the point of interaction. For example, recommendations can be presented to the retail customer who logs in to a retail website.

Use Predictive Customer Intelligence to increase revenue through targeted recommendations

This case study describes how IBM Predictive Customer Intelligence can be used to ensure that an online retail electronics customer progresses through the purchase funnel: from visitor to product viewer, from product viewer to buyer, and from buyer to repeat buyer.

Customer profile

Ross Wilson is a customer who registered for an online retail site recently. Ross has only visited the website a few times and has not spent much time viewing individual products. He has not purchased anything from the website. Ross's lifetime value is Medium, his buying behavior is In-Store Medium Value Offer Seeker, and his demographic cluster is High Income Value. These are predicted values derived from data from both the online store and the physical store.

Stage 1: Visitor

Ross visits the retail website. Based on the demographic profile, visits, and in-store (physical store) purchase details, the Category Affinity model predicts that Ross is most likely to be interested in MP3 players. The website takes Ross to the MP3 player home page. The Inventory-based suggestion model looks at the available stock that is expected to run high on inventory and generates a list of eight MP3 players. The model is based on time series forecasting, current inventory, and the demand variance of all the products. The model also generates a single offer in the MP3 player area for Ross, showing him a clearance item that he currently ignores.

Stage 2: Product viewer

Ross clicks the **Product Details** link for Surf Force MP3 XR32-Silver. The customer segmentation model changes Ross's classification from Visitor to Product Viewer. As a result, two new offers, aimed at progressing Visitors to Product Viewers, are shown to him.

Stage 3: Buyer

Ross decides to purchase this product and adds it to his cart. When Ross completes his purchase of the MP3 player, he is categorized as a Buyer, and is now eligible for a different offer.

The new offer uses purchase information, rather than the browsing information. A market basket recommendation of Wave TV 33 Color is made, based on the past purchase patterns of different customers.

Stage 4: Repeat Buyer

Ross selects the offer, adds the Wave TV to the cart, and checks out. This makes Ross a Repeat Buyer.

Ross receives an offer of 10% extra discount, which is targeted at repeat buyers. Ross also receives a recommendation for a Sonic TLR Thunder Receiver. The choice of this item is based on both of the items that Ross purchased, and is just one of many recommendations. The best recommendation is chosen based on prioritization of offers, generated by an optimization equation that is based on price and cost of the product, the probability that the customer will respond, and expected profit.

Before Ross leaves the site, he decides to look at Big Screen TVs. On an impulse, Ross adds a Wavestation 5000 to the cart.

After reconsidering, Ross removes the TV from the cart, and is offered a discount on the same item.

Because Ross is a Repeat Buyer and he abandoned a cart, he is provided a customized discount. This discount is based on the Price Sensitivity model, which is triggered when the cart is abandoned.

The Price Sensitivity model is developed based on the past purchase behavior of customers of various segments. This model uses two variables to identify the best discount for this customer for this product. The model first takes in the buying behavior segment, which for Ross is In-Store Medium Value Offer Seeker and then the category of product for which discount should be given, "Big Screen TV" (the category of the product he abandoned). Based on both of these factors, the price sensitivity model looks up the best discount that is based on previously developed regression-based models. Ross is given a discount of 5% and the new price is \$1044.99, compared to the recommended retail price of \$1099.

Define the offers to be presented to retail customers

Monica, the marketing manager, defines the company's marketing strategy. She determines which offers might be presented to customers when they visit the online store, based on their customer profile data. Customer profile data includes demographic information, behavioral data, purchase history, and attributes such as segment membership and category affinity, which are derived from predictive models.

Monica's goal is to increase revenue by increasing the rate at which visitors to an online website become product viewers, and viewers become buyers.

Monica looks at the browsing, carting, and purchase metrics for various products. The metrics are presented by using IBM Cognos Business Intelligence reports, which you can drill down on to analyze the online activity pattern of any segment and product combinations.

She sees that there is a low conversion of views into purchase across all product categories and demographic clusters, and decides to make the offers more compelling.

She also sees that the activity is high in the MP3 player section, and decides to apply the Market Basket Analysis model. Market basket analysis allows a retailer

to gain insight into the product sales patterns by analyzing the historical sales records and customer's browsing behavior.

The campaign

In this case study, the example campaign called Predictive Online Promotions is determined by the online purchase funnel that converts Visitors to Product Viewers, Viewers to Buyers, and then Buyers to Repeat Buyers. The online purchase funnel helps you to understand the buying behavior of an average customer so that you can optimize each transaction within a session to move through the funnel.

Using a retail promotions application that is developed in IBM Analytical Decision Management, customers are segmented into categories in real time by a segmentation model. The model predicts whether the customer is a Visitor, Product Viewer, Buyer, or Repeat Buyer, by using data on past activity, and data from the current session. The model is scored in real time as customers move through the store, capturing each new action as input. For example, a customer might progress from Visitor to Viewer to Buyer within a single session, and be presented with appropriate offers at each stage.

A campaign is created in IBM Analytical Decision Management to correspond to each of the four segments defined. Within each campaign, offers are allocated by using segment rules that might also use predictive models. For example, the Clearance item offer is only presented to customers in the Visitor segment with a history of page views for a relevant product line. In this way, a customer's personal preferences can be matched with the items currently on clearance.

Offers can be tailored further by specifying the value of a field or model output to be returned with each offer. For example, for the Clearance offer, the specific product (for example, MP3 Player) that is returned with the offer is selected by the Inventory-Based Suggestion model. This model compares inventory with the customer's known purchase preferences. One model determines which customers are eligible for the offer, and another model determines which product to offer.

Determine the best offers for retail customers by creating business rules

Alex, a business analyst, uses IBM Analytical Decision Management to bring together the company's business rules and predictive models to determine the best possible offer for a customer.

Alex creates business rules under each campaign to determine which actions are valid for any single customer. Here are some examples of what you might set up rules for in the current example, the Predictive Online Promotions campaign.

Visitor

- Clearance item
- Earn double loyalty points

Product viewer

- Value deal
- Popular products
- 1000 loyalty points

Buyer Best in line

Value deal

Repeat buyer

10% additional when you purchase

Special discount for you

For example, the Value deal offer is targeted at customers in both the Product Viewer and Buyer segments who are categorized as offer seekers based on previous buying behavior.

A Market Basket Analysis model is used to identify the next likely purchase for the current customer, and returns this item with the offer. For example if the model returns a Gamma Smartphone, the offer presented to the customer would be “Value deal Gamma Smartphone”.

For the Repeat Buyer campaign, the offer to be returned is determined by a model that returns a recommendation based on the customer’s known sensitivity to price changes.

You can run a simulation for a campaign using sample data to see potential counts of each offer.

Using an optimization equation you can determine the top three offers to be presented to customers based on the following factors:

- Price and cost of the product recommended.
- Probability that the customer will respond.
- Expected profit.

Here is an example of an optimization equation:

$$(\text{PROBABILITY TO RESPOND} * (\text{LIST PRICE} - (\text{UNIT COST} + \text{OFFER VALUE}))) - \text{OFFER COST}$$

For more information, see *IBM Analytical Decision Management Application User’s Guide*.

Build predictive models for retail promotions

David, a data scientist (data engineer), builds predictive models by using IBM SPSS Modeler.

The models are used to gain information for the following purposes:

- Understand the purchase behavior of customers to predict the product lines in which a customer is interested.
- Profile customers based on demographics, online behavior, and buying behavior to help in providing the right offer for the customer at the right time.
- Identify the next product that the customer might be interested in purchasing.
- Determine customer affinity towards product lines by understanding the customer demographic, purchase, and browsing information.
- Find the patterns of customers who are converted to buyers to determine which business rules have the highest impact.
- Determine the extent to which the price of a product affects a customer’s purchase decision.

- Identify the products that have excess inventory so that you can make real-time recommendations to the customers based on the combination of category affinity and the excess inventory.

The scores, segments, and results from these models are used as inputs within IBM Analytical Decision Management by Alex the business analyst as he creates his business rules.

For more information, see “Predictive models in the Retail sample” on page 31.

Use Predictive Customer Intelligence to build cross-channel loyalty: in-store promotions

This example demonstrates how to use IBM Predictive Customer Intelligence to build cross-channel loyalty by using information gained from online activity to determine the best offers to use for customers in a physical store.

Billy Bauer walks into a physical store and purchases a Gamma Smart Phone-2090 Silver. He is also a registered customer in the online website. The web analytics tool recorded that he abandoned a cart to which he added a D2G Performance System (a computer). Billy gets a specialized discount for the computer, that is based on his buying behavior profile, In-Store Medium Value Offer Seeker, and the historical price sensitivity analysis for that cluster on the computers category. Billy is offered a price of \$662.99, which is a 15 percent discount on the retail price of \$779.99. This example highlights how data, which was collected from the customer's online browsing, can be effectively used to give an offer in the physical store, increasing cross-channel loyalty.

Ross Wilson walks into a store and purchases the same Gamma Smart Phone-2090. However, he is given a different offer than Billy Bauer. The offer is personalized, based on Ross's propensity to buy big screen TV products, MP3 players, and smart phones, and the profit that these offers generate based on the margin and propensity to purchase. The offer that Ross receives is for a discounted MP3 player. The MP3 player is part of the Monthly Specials Campaign, which is designed to sell products that are expected to have high inventory.

A third customer, Adrian Alba, who is identified by the Customer Lifetime Value model as a High Value Customer, is provided with 500 Loyalty points from the Rewards Campaign. This is predicted to make him a more loyal buyer.

The campaign

Speakers and receivers have the highest amount of activity associated with them. The marketing manager sees that big screen TV products, smart phones, and MP3 players that have a high inventory are not selling. She decides to design campaigns to sell the big screen TV products, MP3 players, and smart phones, by using a retail promotions application that is developed in IBM Analytical Decision Management.

To clear the excess inventory in big screen TV products, MP3 players, and smart phones, the marketing manager decides to give discount coupons on those items. To select the appropriate people to receive the coupons, she uses the Category Affinity model, which is scored in batch to give out offers to customers whose propensity to purchase a product is more than 30%.

The Rewards Campaign is defined to reward customers based on their value to the retailer. Lifetime Value (LTV), which is an output of the Buying Behavior model, is used to determine whether a customer belongs to the High, Medium, or Low Value group. In each case, the customer is given 500, 200 or 100 loyalty points.

The Customized Pricing offer is used to give attractive discounts on products that the customer has abandoned in the online store. This example highlights how cross channel loyalty can be improved by detecting items that the customer abandoned online and then providing the customer with an offer on those items in the physical store.

The propensity to purchase for big screen TV products, MP3 players, and smart phones, is taken from the Category Affinity model.

The price and cost information for the products is based on actual numbers, while the rewards are based on a prediction of the customer's future purchases.

Insurance case study

By using IBM Predictive Customer Intelligence, insurers can deliver the most appropriate action to each customer during each interaction by analyzing customer data.

Insurance companies have access to large amounts of customer data, and building long term customer relationships by making use of this data is key to their success. The type of data insurers can use is customer profile data, customer policy data, customer claims data, customer complaints data and other transactions, as well as social, and web interaction data.

Use Predictive Customer Intelligence to drive cross-sell opportunities

A multi-line insurance company is facing a problem of declining new premiums from its Life Insurance line of business.

Alex, the business analyst, analyzes the year-to-date data, and identifies that this decline is attributed mainly to the fall in the sale of Term Policies. Further analysis also reveals that there is a large proportion of existing Auto and Home Insurance customers who do not own a Term Policy.

Monica is a marketing manager who is responsible for defining her company's marketing strategy. Monica decides to launch a Cross-Sell campaign to target existing Auto and Homeowner customers for selling Term Policies.

David, the data scientist (or data engineer), builds models by using IBM SPSS Modeler to derive valuable insights from the customer's demographic, transactional, and behavioral data. IBM Predictive Customer Intelligence uses the real-time predictive insights, customer profile data, and current interaction data to recommend the best offer during each customer interaction. IBM Enterprise Marketing Management (IBM EMM) is integrated with the Customer Relationship Management (CRM) software to deliver the campaign offers at the Customer touchpoint.

Customer profile

Grace Wright calls the Insurance company and requests a quotation for a new Life-term policy, with a policy tenure of 25 years. Carolyn, a call center agent, opens her Customer Relationship Management (CRM) system and looks at Grace's profile. She sees that Grace is 32 years old, and is married with one child. Grace is a homeowner, and the household has two insured cars and three policies. Carolyn can also see other information. For example, Carolyn can view complaints that Grace made about the insurance policies. Carolyn can request to see more insights about Grace. IBM Predictive Customer Intelligence uses rich predictive modeling capabilities to view the segment that Grace is in (financially sophisticated), the customer lifetime value (high) and the churn propensity (high). The churn propensity is predicted in real-time based on the latest customer data.

Actions

Carolyn selects the type of interaction (Get Quote) and the product line (policy type of life-term).

Carolyn updates the customer interaction data, and requests a quotation for a policy tenure of 25 years and a coverage amount of \$200,000. IBM Predictive Customer Intelligence provides two offers that are based on customer profile data, predictive insights, and current interaction data:

- Standard premium regardless of health, raised from \$300000 to \$600000.
- Ultra Select Pricing.

Grace then requests a quotation for a policy tenure of 35 years. Using the new interaction data, IBM Predictive Customer Intelligence recommends an extra offer of the last two years of premium waived. Grace accepts the offer and decides to buy the policy. When the offer is accepted, the predictive data is rescored in real time. Grace's churn propensity is reduced.

Define the offers to be presented to insurance customers

Monica is a marketing manager who is responsible for defining her company's marketing strategy. She determines which offers individual customers can receive based on their customer profile data and their real-time interaction data.

Monica determines which offers individual customers should receive based on the customer profile data and the real-time interaction data. This customer profile data can include the following:

- Customer claims data.
- Complaints data.
- Policies data.
- Customer demographic data.
- Predictive Insights such as a customer's segment, and likelihood of churn.
- Real-time data, for example, context of the call – Get a quote for an insurance product, complaint, and so on.

The campaign

Monica creates a campaign that is named Cross-Sell Campaign by using IBM Campaign, part of the IBM Enterprise Marketing Management (EMM) suite. The purpose of the campaign is to cross-sell life insurance policies to existing auto insurance customers.

The Cross-Sell Campaign contains offers to be presented within the insurance call center application. The details of the offer vary based on individual customer characteristics. The results of the offer, whether accepted or rejected, are recorded in the customer's interaction history, and that information can be used for follow-up marketing offers and campaigns.

Customers are segmented into categories in real time that are based on the call center interaction point, such as the context of the call and real time scoring, and the campaign's business rules.

For example, you might have a flowchart with a decision node that assigns customers to five segments:

- High value customers.
- Loyal customers.
- High risk customers.
- Low risk customers.
- High value customers seeking long tenure.

In this example, customers are assigned to the High Value Customers segment if they have a customer lifetime value (CTLV) of high, if their customer segment is Financially Sophisticated, and if they are a home owner.

Customers are segmented based on off-line and real-time customer data, and offers are associated to customer segments.

After each offer is evaluated, there might be more offers eligible than there is time or space to present. For example, a customer is eligible for three offers, but you want to present only the top two offers. In this case, the offers would be presented based on the marketer's score values. The marketer's score can be based on a calculation of an expression by using customer and offer attributes, a real time learning algorithm, or a combination of these methods.

Optimization of offers

You can further optimize offers by using the External Callout connector.

Using the External Callout connector, you can override a single offer's score by calling an IBM SPSS Scoring Service. Also, by using the External Callout connector, you can invoke the IBM SPSS score to present eligible offers.

In the insurance case study, you can invoke the IBM SPSS Scoring Services in real time to get the latest Churn Propensity based on the most current data. The Churn Propensity score is then used as the marketer's score.

For more information, see “The External Callout connector” on page 54.

Determine the best offers for insurance customers by creating business rules

Alex, a business analyst, uses IBM Analytical Decision Management to bring together the company's business rules and predictive models to determine the best possible offer for a customer.

The insurance company has a cross-sell campaign. Alex creates rules to target offers to customers depending on which segment they fall in.

The following examples show offers that are targeted to customers in different segments:

Young family with high churn propensity

Life term insurance with free double indemnity rider

Young family with browsing targeted policies

Life term insurance

Young family with high churn and browsing specific targeted policies

Juvenile life insurance

Divorced and no home ownership

Renter's insurance

Offers are allocated so that the first offer in the list that applies is presented to the customer.

Example of a rule in IBM Analytical Decision Management

The rule "Young family with high churn propensity" takes into consideration the following points:

- Current lifestage is "Young Family".
- Product recently browsed is "Life Term".
- Insurance Policy Recommendation is "Life Term".
- Churn propensity score returned from a real-time model is greater than or equal to 0.7.

Young family with high churn propensity rule

Lifestage Current Segment(Current Segment -) Young Family

PRODUCT_BROWSED_1 = Life term

Insurance_Policy_Recommendation(RECOMMENDATIONS -) =Life Term

Auto Churn(\$RRP-CHURN -Propensity) >=0.7

Build predictive models for an insurance company

David, the data scientist (or data engineer) builds models using IBM SPSS Modeler to achieve the following results:

- Segment customers based on financial sophistication, by using financial and demographic data. In this example, financially sophisticated customers have a relatively high credit score, three types of insurance policies, a relatively high income, and are mostly home owners between 30-35 years of age.
- Predict a customer's propensity to churn, by using information about the customer, such as household and financial data, transactional data, and behavioral data.

- Predict customer lifetime value (CLTV) based on the revenue that the customer brings to the company, the cost of maintaining the policies, the cost of retention, and the likelihood of the customer surrendering the policies in the near future.
- Predict the life policies likely to be bought by customers at different life stages by using historical purchase data. Segment the customers based on their current life stage and predict the life policies most relevant to their life stage.

David uses IBM Social Media Analytics to identify life events by analyzing social media posts to detect life events with potential business value for the Insurance company. Life events such as "New baby born" can act as triggers for cross-selling life policies.

David also uses IBM SPSS Modeler Text Analytics to analyze the customer comments in real time to predict the latest customer sentiment score. IBM SPSS Modeler Text Analytics extracts negative words (sentiment) from customer comments to get a negative sentiment score.

The scores, segments, and results from these models are used as inputs within IBM Analytical Decision Management by Alex the business analyst as he creates his business rules.

Energy and Utilities case study

Energy and Utilities organizations have the challenge of finding, producing, and delivering energy and other natural resources efficiently due to decreasing resources. In a highly competitive industry, Energy and Utility organizations manage high costs and must comply with stringent regulatory policies.

IBM Predictive Customer Intelligence can be used by Energy and Utilities organizations to retain customers and increase customer satisfaction by proactively resolving customer issues. Predictive Customer Intelligence can also help you to identify areas for cross-sell and up-sell across all lines of business.

Retain customers

A call center agent uses the data in the call center application to do the following activities:

1. See the customer's energy usage to verify whether the customer is on the most appropriate contract for his or her needs.
2. See the customer's level of social network influence.
3. See what actions and offers the customer is eligible for.

Determine the best offers for customers by creating business rules

Using IBM Analytical Decision Management, the business analyst creates business rules to determine which actions are valid for a customer. For example, you might create a rule that targets retention actions to customers who are receptive home owners, who have a home over a specific size, and where energy usage exceeds a baseline.

Predict churn, customer satisfaction, and propensity to respond to offers by creating predictive models

Using IBM SPSS Modeler, a data modeler creates predictive models to predict the following factors:

- Customer satisfaction.
- The probability that customers will miss payments.

- The probability that customers will need a demand response program. A demand response program is where customers reduce their energy use at times of peak demand to save money.
- The most profitable Rate Plan.

The results from these models are used in Analytical Decision Management to create business rules.

Banking case study

This case study demonstrates how IBM Predictive Customer Intelligence can be used in the Banking industry.

This case study uses the IBM Enterprise Marketing Management suite along with IBM Predictive Customer Intelligence and a mobile banking application. IBM Interact is used to provide personalized offers. For more information about IBM Enterprise Marketing Management, see Chapter 4, “Integration with the IBM Enterprise Marketing Management suite,” on page 53.

This case study demonstrates how to perform the following actions:

Target the right promotional offers to the right customer

Using IBM Campaign, the marketing manager determines which offers customers will receive based on their customer profile data and their real-time interaction data.

Predict churn, category affinity, and the likelihood of defaulting on credit card debt

Using IBM SPSS Modeler, a data modeler creates predictive models to predict the following factors:

- A customer's propensity to churn.
- What product or service the customer will be most interested in.
- Whether or not a customer will default on credit card debt.

The results from these models can be used to generate scores in the IBM SPSS scoring service, which are then used by IBM Interact to determine which offers are presented to customers.

Use Predictive Customer Intelligence to target promotional offers to the right customer

This case study describes how IBM Predictive Customer Intelligence can be used to target promotional offers to the right mobile banking customer. It demonstrates how IBM Predictive Customer Intelligence can be used to ensure that customers are contacted at the right time, and with the best possible actions.

The mobile banking application connects directly to IBM Enterprise Marketing Management by using the IBM Interact API.

The banking customer

The banking customer is in an airport departure lounge and wants to review her recent account activity before shopping in a duty free store. She logs in to her mobile banking application to display her recent transactions. The customer's recent account activity shows two purchases of online airline tickets. The amounts indicate that these tickets might be for international flights.

IBM Predictive Customer Intelligence offers the customer a discount on travel insurance directly to her phone. The discount is based on the recent purchases of airline tickets.

The customer accepts the offer, and an interactive form appears on her smartphone screen, allowing the customer to enter her details. The customer accepts the offer.

Predictive Customer Intelligence incorporates this new activity from the customer in real time and offers the customer activation of her credit card for international use. The customer decides not to accept the offer. When she rejects the offer, the promotion disappears from the screen of her phone.

Mobile banking workflow

The following steps illustrate how IBM Predictive Customer Intelligence determines which offers are to be presented to the customer.

1. The customer logs in to her mobile banking application through her smartphone to display her recent transactions.
2. The customer's smartphone number and location are passed to IBM Interact (part of the Enterprise Marketing Management suite). IBM Interact, which provides personalized offers and customer profile information in real time, retrieves the customer profile from the IBM Predictive Customer Intelligence database by using the customer ID.
3. An IBM Interact connector, the External Callout connector, calls the IBM SPSS scoring service, passing the customer profile and location data. The IBM SPSS scoring service provides real-time predictive scores or values such as a customer segment name or propensity score.
4. The scoring service sends the score back to IBM Interact, overwriting the Marketer's score to generate an offer that is targeted specifically to the customer.
5. IBM Interact delivers the offer to the customer's smartphone.
6. The customer response from the smartphone is logged in the IBM Interact database along with the offer that was made.

The number of offers that are shown to the customer can be defined by using IBM Interact.

Web service calls are used to connect a mobile banking application with IBM Interact and IBM Predictive Customer Intelligence.

For more information about IBM Enterprise Marketing Management, and IBM Interact, see Chapter 4, "Integration with the IBM Enterprise Marketing Management suite," on page 53.

Define the offers that banking customers can receive

Monica is a marketing manager who is responsible for defining her company's marketing strategy. She determines which offers individual customers can receive based on their customer profile data and their real-time interaction data.

Customer profile data and real-time data

The customer profile data can include the following data: financial status, customer behavioral data, call, and text volume, products that are owned, contract details, predictive model scores, prior transactions, and campaign responses.

Real-time data can be the current location, such as an airport departure lounge, and the action that the customer is taking, such as buying an airline ticket.

The campaign

Monica creates a campaign that is named Mobile Banking Campaign by using IBM Campaign, part of the IBM Enterprise Marketing Management (EMM) suite. The objective of this campaign is to present the right offer to the right customer.

In IBM Campaign, you define offers, and you define which customer segments will be presented with which offer, when a series of conditions are met. For example, are the customers in the affluent segment, and do the customers have a churn score below a certain level?

The marketer's score determines the order in which an offer is presented to the customer. This score can be overwritten by an external callout to the IBM SPSS Scoring service, and this enables you to tailor the offer to the customer's situation.

Build predictive models for a banking customer

David, the data scientist, builds predictive models by using IBM SPSS Modeler.

These models are used to perform the following actions:

- Predict what product or service the customer will be most interested in.
- Predict whether or not the customer is likely to default on their credit card debt.
- Profile customers into groups with similar demand characteristics, such as young educated and middle income, or affluent and middle aged.

This information is used to determine the most appropriate offers to make to the banking customers, and is used in conjunction with IBM Enterprise Marketing Management.

For more information on the predictive models used in the Banking case study, see “Predictive models in the Banking sample” on page 49.

Chapter 3. Predictive models

Use models to predict what is likely to happen in the future, based on patterns in past data.

For example, models can predict the following situations:

- How likely it is that a customer will churn in the next quarter.
- Whether a customer will be a promoter of a service, or a detractor
- How valuable the customer is in terms of future revenue

Models can be used in the same way as business rules. However, while rules might be based on corporate policies, business logic, or other assumptions, models are built on actual observations of past results, and can discover patterns that might not otherwise be apparent. While business rules bring common business logic to applications, models lend insight and predictive power. The ability to combine models and rules is a powerful feature.

Data sources

You or the administrator need to specify the data sources to use in the IBM Predictive Customer Intelligence solution for modeling, analysis, simulation and testing, and scoring.

Use IBM Analytical Decision Management to plan the model and to decide which data sources to use. You need the following types of data in the modeling process:

Historical or analytical data

To build the model, you need information about what to predict. For example, if you want to predict churn, you need information about customers such as their complaints history, number of months since they upgraded their plan, sentiment score, demographic history, and estimated income. This is often referred to as historical data or analytical data, and it must contain some or all of the fields in the project data model, plus an additional field that records the outcome or result that you want to predict. This extra field is used as the target for modeling.

Operational or scoring data

To use the model to predict future results, you need data about the group or population that you are interested in, such as incoming claims. This is often referred to as operational data or scoring data. The project data model is typically based on this data.

You can use the following types of data sources:

- A database that supports ODBC, such as IBM DB2®.
- An Enterprise View that is defined in IBM SPSS Collaboration and Deployment Services.
- A file that is used by IBM SPSS Statistics, such as a text file (txt), or a comma separated file (csv).

When you add a new data source, map all of its fields to ensure compatibility with the project data model. For example, if the project data model requires a field named purchase with values Yes and No for the measurement level flag, then any

data source that you use must have a compatible field. If the field names are not identical, they can be mapped accordingly. Note that the input and associated mapped field must have the same data type.

You can characterize the information that each data field represents. Define a measurement level to determine how a given field is used in business rules, models, or other applications.

You can derive additional fields or attributes for the application by using the expression manager. For example, if you use banking data, you may want to create an expression that shows the ratio between a customer's income and the number of loan accounts that the customer has. Expressions are always numeric with a measurement of Continuous; this cannot be changed.

To enforce corporate-wide policies, use global selections to choose the records to include or exclude from processing by the application. For example, you might have a corporate-wide policy to exclude customers with poor credit or payment histories from future mortgage campaigns. Global selections can be particularly effective when used in combination with shared rules. Shared rules are saved as separate objects that can be used by multiple applications. If the shared rule changes, all applications that use the rule can then be updated.

Data mining that uses IBM SPSS Modeler focuses on the process of running data through a series of nodes. This is referred to as a stream. This series of nodes represents operations to be performed on the data, while links between the nodes indicate the direction of data flow. Typically, you use a data stream to read data into IBM SPSS Modeler, run it through a series of manipulations, and then send it to a destination, such as a table or a viewer.

For example, to open a data source, you add a new field, select records based on values in the new field, and then display the results in a table. In this example, your data stream would consist of the following nodes:

- Variable File node, which reads the data from the data source
- Derive node, which adds the new calculated field to the data set
- Select node, which uses the selection criteria to exclude records from the data stream
- Table node, which displays the results of your manipulations on screen

For more information about these features, see the IBM SPSS Modeler Help (http://www-01.ibm.com/support/knowledgecenter/SS3RA7_16.0.0/com.ibm.spss.modeler.help/clementine/entities/clem_family_overview.htm).

Predictive models in the Telecommunications sample

A number of predictive models are provided in the Telecommunications sample.

To install the samples, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

The following models form the basis of the predictive models in the Telecommunications sample:

Churn model

Customers likely to churn from the current list of active customers can be predicted.

Customer Satisfaction model

Customer satisfaction is determined by the net promoter score.

Association model

Customers can be profiled and assigned to segments.

Response propensity model

You can determine the correct channel to reach the customer, and the probability that the customer will respond.

Predicting churn

Churn is the measurement of subscribers who ended their contract or services. The objective of the Churn Prediction model in the Telecommunication sample is to predict the customers likely to churn from the current list of active customers.

The inputs for the example Churn Prediction model are complaint history, number of months since the customer upgraded the plan, sentiment score, customer demographic history, and estimated income. The example stream for predicting churn is named Churn Prediction.str.




Figure 7. A model that computes churn scores for customers

Data preparation for churn prediction starts with aggregating all available information about the customer. The data that is obtained for predicting the churn is classified in the following categories:

- Transaction and billing data, such as the kind of services subscribed, and average monthly bills.
- Demographic data, such as gender, education, and marital status.
- Behavior data, such as complaints data and price plan migration data.
- Usage data, such as the number of calls and the number of text messages sent.

Data is filtered for modeling in two stages:

1. Data not relevant to some customers.
2. Variables that do not have adequate predictive significance.

A CHAID algorithm is used to predict churn. A CHAID algorithm generates decision trees. A decision tree model is selected over logistic regression because the rules that come out of the decision tree help to understand the root cause of churn better.

The sentiment score is derived from the customer comments text and is an important predictor of churn. Sentiment score considers both the current sentiment score and historical sentiment score.

Other important predictors that are identified during the data understanding and modeling phase are estimated income, number of open complaints, number of closed complaints, time since the last plan upgrade, and the education level of the customer.

Along with the probability of churn occurring, the propensity to churn is calculated by the model. The propensity to churn is widely used in the IBM Analytical Decision Management application. The rule explanation and rule description nodes map the rule identifier number that is generated by the model to the explanation of the rule in English.

Customer satisfaction

Customer satisfaction in the Telecommunications sample is determined by the Net Promoter Score (NPS).

The Net Promoter Score is based on the perspective that every company's customers can be divided into three categories:

- Promoters are loyal enthusiasts who keep buying from a company and urge their friends to do the same.
- Passives are satisfied but unenthusiastic customers who can be easily wooed by the competition.
- Detractors are unhappy customers who are trapped in a bad relationship with the company.

The Net Promoter Score is obtained by asking a set of customers a single question: "How likely is it that you would recommend our company to a friend or colleague?" Customers are asked to answer on a 0 - 10 rating scale. Based on the score that they provide, they are categorized as Promoter (if the score is 9 or 10), Passive (if the score is 7 or 8), or Detractor (if the score is 6 or less).

The objective of the Net Promoter Score model is to identify the distinguishing characteristics of the customers who fall into the three categories. The net promoter score model is then used to predict which category a customer would fall into, without asking the question "How likely is it that you would recommend our company to a friend or colleague?" This model helps to dynamically track the change in the Net Promoter Score of a customer.

The example stream for identifying the Net Promoter Score is named `Satisfaction.str`.

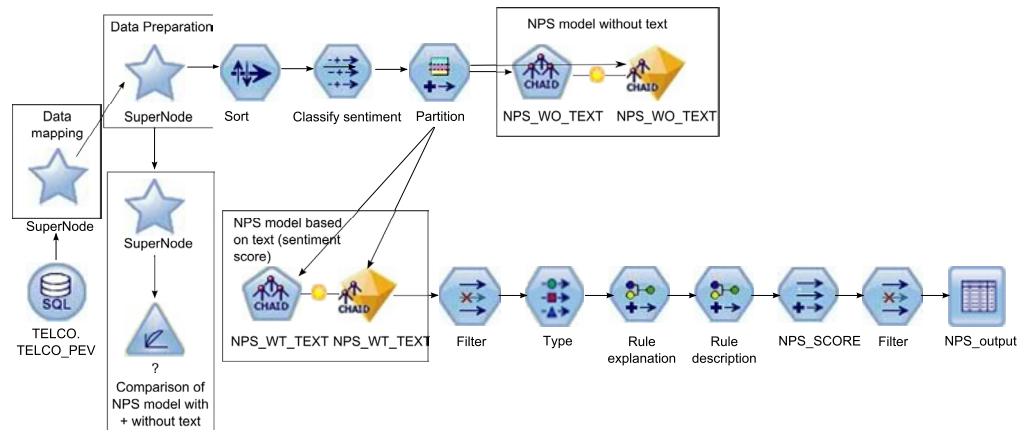


Figure 8. Stream for identifying customer satisfaction

Historical data comes from a sample of customers who answered the question. Customers for whom there is no score are considered to be operational data, whose satisfaction group needs to be predicted for the first time. The Customer Satisfaction model can be used to predict scores for customers who do not have a net promoter score.

The sentiment score, along with the number of open complaints, employment status, and estimated income, are identified to be the key variables that affect the prediction of satisfaction group. The sentiment score is focused on capturing the negative sentiments across various attributes, such as network and service. A sentiment score of zero means that the customer has not expressed any negative sentiment. A sentiment score of two means that the customer has expressed negative sentiment in two predefined categories. Six categories were defined, and so the maximum sentiment score is 6.

The sentiment score that is used in the example database is an average value of the most recent sentiment score calculated and the previous sentiment score of the same customer. Where a customer expressed negative sentiment on a single category, and then expressed multiple positive comments, the sentiment score would be mildly negative, although close to zero. For the purposes of satisfaction modeling, to avoid categorizing the customer as mildly negative, sentiment scores less than 0.6 are rounded to zero.

Assigning offers

An Association model is used to assign the right offer to a customer. It uses the customer's segment (for example, Platinum) and predicted net promoter score group (for example, Promoter) to determine an offer (for example Phone Plan).

Segmentation is the process of profiling customers into groups with similar demand characteristics. The example stream for profiling customers is named `AssociationModel.str`.

The following shows an example association model.

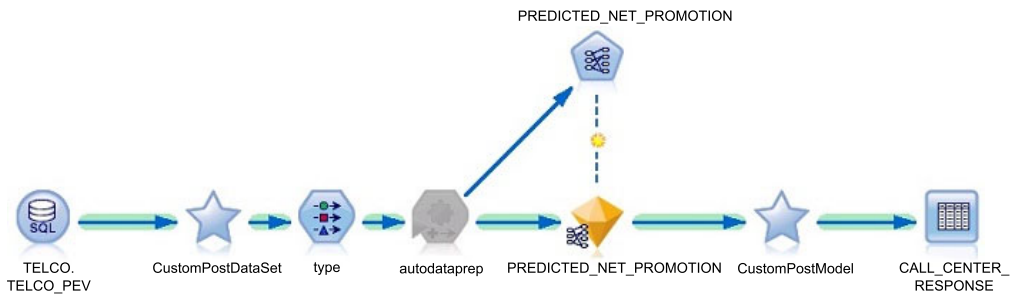


Figure 9. Association model for the Telecommunications case study

Targeting offers to customers with the response propensity model

It is important to target offers to the correct customers, through the correct channel.

The Response Propensity model determines the correct channel to reach the customer, and determines the probability that the customer will respond.

The example stream for determining response propensity is named ResponsePropensity.str.

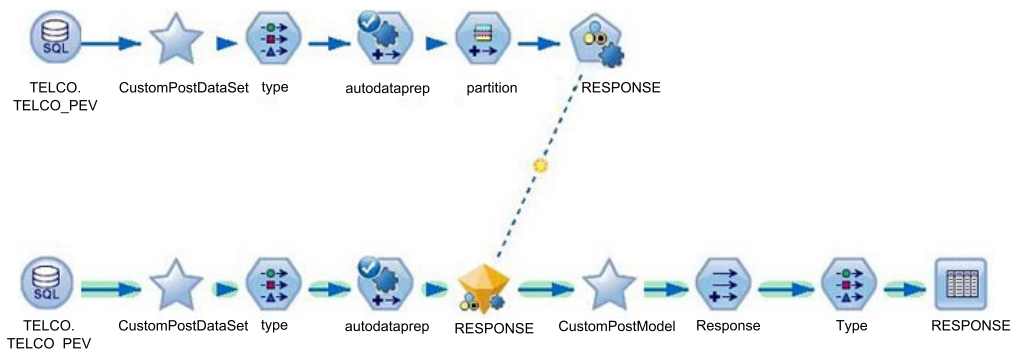


Figure 10. Response propensity model

You can use the results of this model to target customers who are likely to respond because they are above a certain threshold, or ignore customers who are likely to result in a minimum profit.

The input for the model is customer demographic information, billing history, customer lifetime value, churn score, net promoter score, and tenure.

The customer's previous offer response data can be used as the input for the current model. The historical data on which interaction points the customer has responded to an offer is taken and the model is trained based on that data.

Telecommunications models in Analytical Decision Management

In IBM Analytical Decision Management, you can combine predictive models with rules to allocate offers in accordance with business goals. You do this by combining selection and allocation rules that are based on the output from predictive models.

There are two main steps:

- Define and allocate offers to determine which offers a customer is eligible for.
- Prioritize offers to determine which offers a customer receives.

Predictive models in the Retail sample

A number of predictive models are provided in the Retail sample.

To install the samples, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

The following models form the basis of the predictive models in the Retail sample.

Customer Segmentation model

Customer segmentation involves profiling customers through demographic segmentation, online behavior segmentation, and buying behavior segmentation.

Market Basket Analysis model

Market basket analysis allows retailers to gain insight into the product sales patterns by analyzing historical sales records and customers' online browsing behavior.

Customer Affinity model

You can determine customer affinity towards product lines by understanding the customer demographic information, purchase information, and browsing information.

Response Log Analysis model

Response log analysis captures the response of customers compared with the recommendations that come from IBM Analytical Decision Management business rules.

Price Sensitivity model

Price sensitivity is the extent to which the price of a product affects a customer's purchase decision. The degree of price sensitivity varies from customer to customer, and from product to product.

Inventory-Based Suggestion model

The Inventory-Based Suggestion model identifies the products that have excess inventory and then makes real-time recommendations to the customers based on the combination of category affinity and excess inventory.

Prepare data for a retail solution

Before data can be used in predictive models, process the online browsing data by using IBM SPSS Modeler streams.

Pre-processing online data

Use a data pre-processing stream to process the browsing behavior data from an online system and load it into a database table in a format that is suitable for modeling analysis. Processing can be done in batch processing mode.

The example stream for pre-processing data is named `Preprocessing_Data>Loading_Stream.str`.

To obtain the online browsing behavior data, deploy a web analytics solution. Web analytics tools allow the flexibility to export the browsing behavior data in different formats such as a comma-separated file, which can then be consumed in IBM SPSS Modeler streams for further analysis. Data might include information such as products browsed, products put into a cart, products abandoned, products purchased, pages viewed by customers, and the product category.

Category affinity target determination

Affinity analysis is a data analysis and data mining technique that discovers co-occurrence relationships among activities that are performed by the customer to understand the purchase behavior of customers. The customer can have both an online and a physical presence. During online shopping, customers can browse, search, and view pages for different products before they make a purchase. The purpose of the Category Affinity model is to get information about product lines in which a customer is interested by understanding the online and buying behavior.

The stream for category affinity target determination that is included in the samples is named `Category_Affinity_Target_Determination.str`.

Use a category affinity target determination stream to process customer online and in-store historical transactional data to find out the following information:

- Products browsed
- Products purchased
- Products abandoned
- Products that are put in a cart
- Onsite searches
- Page views

The stream should process each activity separately so that you can prioritize the order of activity as purchase, browse, search, and page views.

Data should be processed in two steps:

1. Aggregation of the data should be done at the product line level so that you can get the number of items per product line for each customer activity.
2. Identify whether a customer has an affinity for a particular product line by comparing:
 - the number of items that are purchased in a category, with
 - the average number of items that are purchased by the total population.

Some product lines are more likely to be purchased in large quantities while some other products would be purchased in small quantities. For example, a customer might buy many writeable DVDs but they might buy a computer once in three years. If you take the number of items or the value of items, the Category Affinity model would show a bias to a select few products.

Customers are likely to search products with different names, and might use related keywords or visit related pages. The pre-processing stream can process available data to get the corresponding product line information and derive the volume of items for each product line under this category.

When a customer searches a product with an exact product line name, a weight of one is assigned to that product line. However, when a customer searches for a super category, the number of product lines in that super category is determined

and all the product lines in that super category are assigned equal weight. For example, a customer searches by the category named Consumer Electronics. It is not possible to know which product line the customer searched for because Consumer Electronics contains three product lines: Computers, MP3, and Smartphone. In such cases, all three product lines have a weight of $1/3 = 0.33333333$.

You can get more insight into customers from their purchase behavior than from their browsing behavior. Consider browsing behavior only when there is no purchase information. Likewise, if there is no browsing behavior, consider searching behavior, and if there is no searching behavior, consider page view information.

The formula for this example is shown:

```
if not(@NULL(Line_Purchased)) then Line_Purchased
elseif not(@NULL(Line_Browsed)) then Line_Browsed
elseif not(@NULL(Line_Search)) then Line_Search
elseif not(@NULL(Line_PageView)) then PageView
else "Default"
endif
```

In the samples, this derived information is stored in the intermediate database table: `RETAIL.CATEGORY_AFFINITY_OUTPUT`. This information is used as input into the Logistic Progression model in the `Category_Affinity_MBA_Segmentation.str` stream.

Defining customer segments

When you define customer segments, you profile the customers into groups that have similar demand characteristics.

You profile customers based on demographics, online behavior, and buying behavior to help to provide the right offer for the customer at the right time.

The example stream for defining customer segments in the Retail sample is named `Customer_Segmentation.str`.

Demographic segmentation

Demographic segmentation is based on age, gender, marital status, number of members in household, education, profession, and income. To get meaningful segments out of the data, missing information is derived based on the other variables. Continuous variables such as age and income are tiled into smaller number of groups.

When there are multiple variables to consider, clustering can be challenging. K-Means clustering can be used to cluster the customer base into distinct groups. Instead of trying to predict an outcome, K-Means clustering tries to uncover patterns from the demographic data that is provided as input. The clusters are formed in such a way that each customer within the segments is similar and different from customers present in other segments. Multiple iterations of K-Means clustering on the customer base are performed to arrive at six clusters to be used for targeting campaigns. Education, income, and marital status are the top three variables that determine the cluster in which a customer is categorized.

Online behavior segmentation

Before online behavior segmentation can begin, the data must be prepared. Data preparation is described in “Prepare data for a retail solution” on page 31. This data is then aggregated to reveal trends in a particular session, and across sessions for an individual customer, and the total population.

The online browsing history is taken, and based on the aggregate information that is prepared in the data preparation stage, a TwoStep cluster model is developed. The TwoStep clustering method is used because it can handle mixed field types and is able to handle large data sets efficiently. It can also test several cluster solutions and choose the best, so the number of clusters does not have to be set at the beginning of the process. Not having to set the number of clusters is an important feature because online data is dynamic in nature, and the algorithm that is used must be able to identify new clusters of users when they form. The TwoStep clustering method can be set to exclude unusual cases that can contaminate results. The clusters that are formed by the TwoStep clustering method depict the qualities of customers in the various stages of the online purchase funnel.

Buying behavior segmentation

The purchase history of the customer, both online and in-store, is collected. Extra details that relate to the products purchased are also collected. The extra details might include: was the item purchased when there was an offer? What was the discount? What is the margin from selling the product?. This information is used to derive multiple variables such as the average purchase value and average number of items that are purchased, across both channels and during discount and regular shopping.

The past purchase behavior helps to predict how likely a customer is to react to various kinds of offers. The past purchase behavior also helps to predict the kind of campaign that would be more suitable for the customer. Based on the online and in-store purchases that are made by the customers, and whether the purchases were made during an offer or at another time, a K-means clustering model is used to derive the various segments. The model identifies two clusters for primarily online customers, and three clusters for primarily in-store customers. For both online and in-store clusters, one segment is identified as offer hunters. Online offer hunters typically purchase high value items and in-store offer hunters typically purchase low value items. Non-offer hunters typically purchase high value items.

Market Basket Analysis

Market Basket Analysis allows retailers to gain insight into the product sales patterns by analyzing historical sales records and customers' online browsing behavior.

Market Basket Analysis is used to increase marketing effectiveness and to improve cross-sell and up-sell opportunities by making the right offer to the right customer. For a retailer, good promotions translate into increased revenue and profits. The objectives of the market basket analysis models are to identify the next product that the customer might be interested to purchase or to browse.

Market Basket Analysis works on relating products in the historical transactions of a retailer. Association rules are generated by using the frequency of a particular item along with the combination of items. Rules with higher lift, confidence, and support are selected for deployment.

The example streams for Market Basket Analysis have the following names:

- `Market_Basket_Analysis_Product_Recommendations.str`
- `Market_Basket_Analysis_Export_to_Table.str`
- `Market_Basket_Analysis Product Lines.str`

The inputs into the streams are the product and categories that are purchased and browsed. The output is used in IBM Analytical Decision Manager to provide an appropriate offer for a customer.

Market Basket Analysis for an online retailer requires two types of data: sales transaction data, and the customer's online browsing behavior data.

Online data that was prepared in the data processing step, provides the information on online purchase and browsing. For more information, see "Prepare data for a retail solution" on page 31. The in-store transaction tables provide the information on physical store purchase.

An Apriori algorithm is used to find association rules between different product categories. Market Basket Analysis is done separately to find associations between products that are browsed, products purchased online, and products purchased in-store. Browsing behavior data is aggregated to get all the products purchased by a customer and all the products browsed by a customer. The Apriori algorithm is then applied on the aggregated data to find the association between different product categories for the products that are purchased and the products that are browsed. For more information about the Apriori algorithm, see the IBM SPSS Modeler User's Guide (http://www-01.ibm.com/support/knowledgecenter/SS3RA7_16.0.0/com.ibm.spss.modeler.help/clementine/understanding_modeltypes.htm?lang=en).

Determine Customer Affinity

You can determine customer affinity towards product lines by understanding the customer demographic information, purchase information, and browsing information.

The example stream for determining product line affinity and product selection for customers in the Retail sample is named `Category_Affinity_MBA_Segmentation.str`.

The input for the Category Affinity Model is the output of the Category Affinity Target Determination model, the Customer Segmentation model, and online transactional data. The outputs from the models are stored in the intermediate database tables, and contain information about the product lines that customers are most interested in, the customer segments, and the customer market basket. This input is used by a logistic regression algorithm. Logistic regression is a statistical technique for classifying records that are based on values of input fields. A multinomial model is used because the target has multiple product lines. The model output is stored in the `CUSTOMER_SUMMARY_DATA` database table and is used in the IBM Analytical Decision Management application.

Response log analysis

A response log captures the response of customers compared with the recommendations that come from IBM Analytical Decision Management business rules.

A response log gives a view of how many recommendations are taken up by customers in the form of offers that are accepted or offers that are declined. This information is logged in IBM SPSS Collaboration and Deployment Services systems tables in the form of XML response logs. The objective of the Response Log Analysis model is to discover the following information:

- The patterns of customers who are converted to Buyers.
- Which IBM Analytical Decision Management rules triggered high conversion rates to determine high impact business rules.

The example stream for the example Response Log Analysis model is named `ResponseLog_Model.str`.

The inputs are the product feedback log, and the response log.

The following models are used:

- Self-Learning Response Model (SLRM) algorithm
- Bayesian Network algorithm

Response log data is captured by the Response service. The Response service logs all the customer responses to IBM SPSS Collaboration and Deployment Services system tables in the form of XML tags. The log contains customer responses such as offers accepted, offers made, customer demographics, actual profit, IBM Analytical Decision Management rules for the offers that are accepted, and other metadata.

A similar method can be used to log product feedback that is given by customers in text format against each product line. This data is queried by XQuery, a query and functional programming language that queries collections of XML data. The data is then loaded into a view and used as a source of data for modeling.

The Self-Learning Response Model (SLRM) algorithm is used to predict the best offers to customers by using past responses to recommendations, and customer demographics. By using the SLRM node, you can build a model that continually updates, or re-estimates, as a data set grows without having to rebuild the model using the complete data set. This model predicts which offers are most appropriate for customers and predicts the probability of offers being accepted. The model predicts the best three offers for a customer. The model also analyses IBM Analytical Decision Management rules to determine which are the most effective rules.

Inventory-based suggestion

Retailers often have the problem of excess inventory, where the value of the inventory depreciates rapidly due to the products becoming outdated. To prevent this problem, retailers use offers in order to clear the excess inventory. The Inventory-Based Suggestion model identifies the products that have excess inventory and then makes real-time recommendations to the customers based on the combination of category affinity and excess inventory.

The example stream for the Inventory-Based Suggestion model is named `Inventory_Based_Suggestion.str`.

The input to the model is transaction data from the online and physical store, and product details, including the current inventory. The output is the product name, and price and cost data to be made available to the customer.

The modeling techniques that are used are time series modeling, and inventory analysis by using IBM SPSS Modeler derive node functions.

Forecasting inventory

You can predict the demand for products a week in advance.

Purchases that are made by the customers both in-store and online are aggregated by day, giving the information on demand for products daily. This information is used as an input for time series modeling. In the Time Series model, the expert option is selected so that IBM SPSS Modeler selects the best fitting model for each of the products depending on their individual characteristics.

Inventory cost analysis

You can calculate excess inventory and the holding cost. Excess inventory is determined by using the current stock, forecasted demand for the period in consideration, and the variation in demand.

$$\text{Excess Inventory} = \text{Current Stock} - \text{Forecasted Demand} - \text{Z score required for the service level} * \text{Expected Variance in demand}$$

Excess inventory is calculated a week in advance. The current stock is taken from the `PACK SIZE` variable in the product table. The forecasted demand for the next seven days is given out by the time series model. The standard deviation for a single day is calculated by using the aggregate node, which gives the values for standard deviation for all products. To get the variance over the length (seven days) the variance is multiplied by $\text{SQRT}(7)$. The variance over a longer duration is an addition of the variance expected for each day, and standard deviation is the square root of variance. The holding cost is then taken to be 25% of the cost of all the products that are in excess inventory.

Real-time recommendations

When a customer's data is entered in IBM Analytical Decision Manager, the product line to which the customer has the highest affinity is selected in real time. If there is a product in that product line that has excess inventory, that product is recommended to the customer. If the customer's affinity is with a product line that does not have excess inventory, then the default product, which has the highest holding cost across all product lines, is recommended to the customer.

Deployment of models in the Retail sample

Scoring streams are based on the predictive models that are described in the Retail case study.

The scoring streams use the same input as the IBM Analytical Decision Management application.

In the Retail sample, there are three scoring streams:

- Online behavior segmentation.
- Browsing Market Basket Analysis at a product level by customer.
- Buying Market Basket Analysis at a product level by customer.

To perform Market Basket Analysis, you must look up the price information in the product table or the product that comes out of the Market Basket Analysis recommendation. In a real-time scoring service, this is not possible, so instead, scripts are used. Information, such as price and cost, is populated into look-up tables by using these scripts. There are two ways to execute the script.

- To execute the script interactively in IBM SPSS Modeler, select **Tool > Stream properties > Execution**. The example script runs the product table output node that gives information on the product price, cost, and full item name. The script then deletes the existing lookup tables and creates new lookup tables that are based on the product table output.
- The second method of execution is to create a job and invoke the script by using the job URI through a web service invocation.

Example script to populate lookup tables

The following script provides an example of how to populate lookup tables. Modify this script for your environment.

```
#####
# Variable Definition
#####
var last_row
var i
#####
delete Price:reclassifynode
execute PRODUCT_TABLE:tablenode
set last_row = PRODUCT_TABLE:tablenode.output.row_count

create reclassifynode connected between Type_Before and Cost:reclassifynode
set Reclassify:reclassifynode.new_name = "Price"
set Price:reclassifynode.mode = Single
set Price:reclassifynode.replace_field = false
set Price:reclassifynode.field = "$A-30 fields-1"
set Price:reclassifynode.use_default = True
for i from 1 to ^last_row
  set "$A-30 fields-1" = value PRODUCT_TABLE:tablenode.output at ^i 1
  set Price = value PRODUCT_TABLE:tablenode.output at ^i 4
  set Price:reclassifynode.reclassify.^"$A-30 fields-1" = ^Price
endfor

#####
delete Cost:reclassifynode
execute PRODUCT_TABLE:tablenode
set last_row = PRODUCT_TABLE:tablenode.output.row_count

create reclassifynode connected between Price:reclassifynode and
  Full_Item_Name:reclassifynode
set Reclassify:reclassifynode.new_name = "Cost"
set Cost:reclassifynode.mode = Single
set Cost:reclassifynode.replace_field = false
set Cost:reclassifynode.field = "$A-30 fields-1"
set Cost:reclassifynode.use_default = True
for i from 1 to ^last_row
  set "$A-30 fields-1" = value PRODUCT_TABLE:tablenode.output at ^i 1
  set Cost = value PRODUCT_TABLE:tablenode.output at ^i 3
  set Cost:reclassifynode.reclassify.^"$A-30 fields-1" = ^Cost
endfor
```

```
#####
delete Full_Item_Name:reclassifynode
execute PRODUCT_TABLE:tablenode
set last_row = PRODUCT_TABLE:tablenode.output.row_count

create reclassifynode connected between Cost:reclassifynode and
  Full_Item_Name_With_Price
set Reclassify:reclassifynode.new_name = "Full_Item_Name"
set Full_Item_Name:reclassifynode.mode = Single
set Full_Item_Name:reclassifynode.replace_field = false
set Full_Item_Name:reclassifynode.field = "$A-30 fields-1"
set Full_Item_Name:reclassifynode.use_default = True
for i from 1 to ^last_row
  set "$A-30 fields-1" = value PRODUCT_TABLE:tablenode.output at ^i 1
  set Full_Item_Name = value PRODUCT_TABLE:tablenode.output at ^i 2
  set Full_Item_Name:reclassifynode.reclassify.^"$A-30 fields-1" = ^Full_Item_Name
endfor

#####
```

Use Retail case study models in IBM Analytical Decision Management

In IBM Analytical Decision Management, you can combine predictive models with rules to allocate offers in accordance with business goals. You do this by combining selection and allocation rules that can be based on simple attributes such as age or job status, or based on the output from predictive models.

The Analytical Decision Management application for the Retail sample is named **Retail Promotions**, and is designed for campaign management in a retail scenario. It is shared by two Analytical Decision models, one model for online promotions, and one model for in-store promotions.

Online promotions

For online promotions, campaigns are designed to target specific customer segments.

The data that is used by the application includes demographic, behavioral data, and purchase history, and attributes such as segment membership and category affinity that are derived from predictive models.

In-store promotions

For in-store promotions, campaigns are organized around business objectives, such as reducing inventory through special promotions or rewarding the most loyal customers.

Input to the Analytical Decision Management application

The IBM Analytical Decision Management application uses the output of the following predictive models as input:

Category Affinity model output

The probability that the customer likes a particular product line.

Segmentation model output

The demographic, online behavior, and buying behavior segment outputs.

Market Basket Analysis output

The output of market basket analysis that is based on browsing and purchase.

Predictive models in the Insurance sample

A number of predictive models are used in the Insurance sample.

To install the samples, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

The following models form the basis of the predictive models in the Insurance sample:

Segmentation model

Customers are segmented based on their financial sophistication. This model enables the insurer to sell insurance policies that are appropriate to the customer.

Churn Prediction model

Predicts a customer's propensity to churn by using information about the customer such as household and financial data, transactional data, and behavioral data.

Customer Lifetime Value model (CLTV)

Predicts customer lifetime value. CLTV is based on the revenue that the customer brings in to the company, the cost of maintaining the policies, the cost of retention, and the likelihood of the customer surrendering the policies soon.

Campaign Response model

Predicts the probability that customers will respond to targeted offers so that you send out offers only to the customers whose propensity to respond is above a particular threshold.

Auto Churn model

Predicts the customers likely to churn from the current list of active customers. This model considers the Auto policy customers only.

Life Stage Segment model

Groups customers based on their current life stage, which would help in recommending the right insurance policy to the customer based on their current life stage segment.

Buying Propensity model

Identifies the life policies that are mostly bought by customers belonging to each life stage segment as defined in the life stage segment model.

Data Processing Stream

Transforms and aggregates data obtained from the customer's visit to the insurer's website so that it can be used to define rules for recommending the right insurance policy to each customer.

Insurance Policy Recommendation model

Recommends the correct insurance policy by considering information such as the customer's web activity data and also the life insurance policy buying propensity for the given customer's life stage segment, based on historical data.

Social Media Analytics (SMA) model

Extracts customer life stage event information from customers social media posts.

Sentiment Scoring model

Extracts the sentiment score from customer comments that are captured while recording customer complaints.

Data used in the Insurance sample

In the Insurance sample, the Insurance company runs a multi-line business.

The following types of data are used.

Customer master data

This includes customer's demographic data, employment and income data, and information about the household as well. POLICYHOLDER and HOUSEHOLD tables capture most of this data. Typically, Master Data Management systems are the source of customer master data.

Customer policy data

This includes aggregated customer information, such as the number and types of policies owned by the customer, total premium being paid by the customer, average claim amount, tenure of the customer, number of complaints, number of claims, and customer sentiment data. POLICYHOLDER_FACT and POLICY_FACT tables capture most of this data.

Customer transaction data

This includes data about all the customer transactions such as the policies purchased, their inception and maturity/renewal date, data related to all of the complaints made by the customer in the past, and also data related to all of the claims made by the customer. POLICIES, CLAIMS, COMPLAINTS, COMPLAINT_DETAILS tables contain this data.

Customer Social Media Data

Apart from the customer data that is available within the enterprise, insurance organizations may also want to get insights from external sources of data. For example, the social media channels where customers post comments about their experiences with their insurers, as well as about their needs and life-events that can potentially lead to an opportunity to sell appropriate insurance products. SMA_DATA and SMA_DATA_ANALYSIS tables captures such external data, as well as the summarized analysis of this social media data.

Note: The fetching of social media data is outside the scope of this solution.

Customer web browsing data

Many insurance organizations today allow their customers to buy or explore their insurance products online through their websites. Technology makes it possible to track customers' activities on their websites, giving them vital insights about the customers' current interest in specific insurance products. Web Analytics tools can be used to analyze customers' website activities and use this information along with other customer data to make the right recommendations to the customers at the right time. ACTIVITY_FEED_DATA, ONLINE_BROWSING_HISTORY and ONLINE_BROWSING_SUMMARY tables contain customers web activity data.

Define customer segments in the Insurance sample

When you define customer segments, you profile the customers into groups that have similar demand characteristics. In the Insurance sample, customers are profiled based on their financial sophistication.

Customers are segmented into financially sophisticated, and novice categories. This means that insurers can target each segment with cross-selling insurance policies that are appropriate to increase the effectiveness of cross-sell campaigns.

The example stream for defining customer segments is named `Segmentation.str`.

The model used is a `TwoStep` cluster.

The inputs for the example `Segmentation` model are customer master data and customer policy data, specifically:

- Demographic data: age, gender, marital status, and employment status.
- Insurance policy related data: insurance lines, policies, premiums, tenure, and insurance score.
- Financial data: income, retirement plans, home ownership status, vehicle ownership.

These inputs are aggregated. As each record is read, based on a distance criterion, the `TwoStep` cluster algorithm decides whether it should be merged with an existing cluster, or used to start a new cluster. If the number of clusters becomes too large (you can set the maximum number), then the distance criterion is increased, and clusters whose distance is now less than the modified distance criterion are merged. In this way, the records are clustered into a preliminary set of clusters in a single pass of the data.

Predict churn in the Insurance case study

The Churn prediction model predicts a customer's propensity to churn by using information about the customer such as household and financial data, transactional data, and behavioral data.

The inputs for the Churn prediction model are customer demographic data, insurance policies, premiums, tenure, claims, complaints, and the sentiment score from past surveys.

The example stream for predicting churn is named `Churn.str`.

Data preparation for churn prediction starts with aggregating all available information about the customer. The data that is obtained for predicting the churn is classified in the following categories:

- Demographic data, such as age, gender, education, marital status, employment status, income, home ownership status, and retirement plan.
- Policy-related data, such as insurance lines, number of policies in the household, household tenure, premium, disposable income, and insured cars.
- Claims, such as claim settlement duration, number of claims that are filed and denied.
- Complaints, such as number of open and closed complaints.
- Survey sentiment data. Sentiment scores from past surveys are captured in the latest, and average note attitude score fields. The note attitude score is derived

from customer negative feedback only. If the note attitude is zero, the customer is more satisfied while as the number increases, satisfaction level decreases.

A CHAID algorithm is used to predict churn. The CHAID (Chi-squared Automatic Interaction Detection) is a classification method that builds decision trees by using chi-square statistics to work out the best places to make the splits in the decision tree. The CHAID model output is decision rules/tree and a predictor importance plot. This output shows you which the most important predictors are for whether a customer is likely to churn or not. For example, the most important predictors could be HOUSEHOLD_TENURE, LATEST_NOTE_ATTITUDE, and NUMBER_OF_POLICIES_IN_HOUSEHOLD.

Understand Customer Lifetime Value (CLTV)

The Insurance sample uses Customer Lifetime Value (CLTV) to understand customer profitability.

CLTV is a commonly used approach to determine how much each customer is worth in monetary terms. It helps insurers to determine how much money must be spent to acquire or retain a customer. CLTV is defined as the expected net profit a customer will contribute to the business over time. Advanced analytics provide new insights on how customer lifetime value can be calculated.

The example stream is named CLTV.str.

The inputs for this model are customer demographic data, policies, premiums, tenure, policy maintenance cost, complaints, and customer survey sentiment.

CLTV is determined by the margin amount that the customer contributes each month and the probability that the customer might churn in any month. Customers with higher margin amount and a lower probability to churn have a high CLTV.

CLTV is derived by the following formula:

$$\sum_{i=0}^N \frac{\text{NetProfit} * C_i}{(1+d)^t}$$

C_i = probability for customer i to generate revenue in time t

N = total number of periods

d = monthly discount rate

t = time of cash flow

The probability C_i can be estimated by using Cox regression. Cox regression is a method for investigating the effect of several variables upon the time a specified event takes to happen. In the context of an outcome such as churn, this is known as Cox regression for survival analysis.

CLTV is calculated by considering the following:

- When the Cox model is scored, the customer's past 'survival' time is considered, and the churn probability is predicted for one to five years.

- The NetProfit value is derived by using following expression.

$$\text{NET_PROFIT} = (\text{TOTAL_PREMIUM} - \text{MAINTENANCE_COST}) * 12$$

- The customer lifetime value is derived as follows:

$$\begin{aligned} \text{CLTV} = & (\text{NET_PROFIT} * C_1 / (1 + 0.12)) \\ & + (\text{NET_PROFIT} * C_2 / (1 + 0.11) ** 2) \\ & + (\text{NET_PROFIT} * C_3 / (1 + 0.1) ** 3) \\ & + (\text{NET_PROFIT} * C_4 / (1 + 0.09) ** 4) \\ & + (\text{NET_PROFIT} * C_5 / (1 + 0.08) ** 5) \\ & + (\text{NET_PROFIT} * \text{POLICYHOLDER_TENURE}) \end{aligned}$$

Where $C_{i=(1,2,3,4,5)}$ is the renewal probability, which is the future value that a customer can bring. The last item is the historical/current value of one customer.

- The CLTV values are further classified into Low, Medium, and High categories by using the following calculations:

```
CLTV_CAT =
  if CLTV <= 30083.625 then 'LOW'
  elseif CLTV > 30083.625 and CLTV <= 46488.000000000007 then 'MEDIUM'
  elseif CLTV > 46488.000000000007 then 'HIGH'
  else 'LOW'
endif
```

Predict customer response to a campaign

Targeting offers to the correct customers is an important part of promotion planning and campaign design. The Insurance sample uses the Campaign Response model to predict the probability that the customer will respond to targeted offers.

The Campaign Response model helps in sending out offers only to the customers whose propensity to respond is above a particular threshold.

The example stream for predicting customer response to a campaign is named Campaign Response Model.str.

The customer's previous offer response data is the input for the model, and the model is trained based on that data.

The Decision List algorithm is used to identify the characteristics of customers who are most likely to respond favorably based on previous campaigns. The model generates rules that indicate a higher or lower likelihood of a binary (1 or 0) outcome. The Campaign Response model considers only those customers who are currently with the insurance company, and not those customers who churned.

Predict churn for auto policy holders

The Insurance sample uses the Auto Churn prediction model to predict the customers likely to churn from the current list of active customers that hold auto policies.

The example stream for predicting the customers likely to churn is named Auto Churn Model.str.

The CHAID decision tree algorithm is used to predict churn. This model is similar to the Churn model, except that this model only considers auto policy data for predicting churn. For more information, see "Predict churn in the Insurance case study" on page 42.

In the example provided with IBM Predictive Customer Intelligence, the most important predictors of churn are the last survey sentiment score (LATEST_NOTE_ATTITUDE), and CLAIM_SETTLEMENT_DURATION.

Group customers based on their current life stage

The Insurance sample uses the Lifestage Segment model to group customers based on their current life stage.

The example stream for grouping customers into their current life stage is named `Lifestage Current Segment.str`.

The model uses simple rules to get the current life stage segment of customer. Some examples of defined segments are:

- Newly married.
- Young family.
- Young and affluent.
- Single.
- Divorced.

Identify the life policies bought by life stage segments

The Insurance sample categorizes customer buying propensity for life policies into life stage segments using the Apriori model.

The example stream is named `Buying Propensity_Model.str`.

The Apriori model is an association algorithm that extracts association rules from data. The algorithm uses the past insurance policy purchase data and provides the buying propensity score for each policy at customer life stage segment level. The output is processed further into a summary format, which is then used by IBM Analytical Decision Management to deliver appropriate offers to customers.

Recommend the correct insurance policy

The Insurance sample uses the Insurance Policy Recommendation model to recommend the correct insurance policy for customers.

The example stream is `Insurance Policy Recommendation.str`.

The Insurance Policy Recommendation model compares insurance policies that are browsed by the customer with the list of insurance policies that have a higher buying propensity for the customer's life stage segment. The model recommends the correct insurance policy according to this data.

Transform and aggregate customer data

The Insurance sample uses the `Data Processing Stream.str` to transform and aggregate data about customer activity on the insurer's website.

To obtain the online browsing behavior data, deploy a web analytics solution.

The stream pre-processes the activity feed data from a web analytics tool and loads it into tables in a format that makes it easy for analysis. The transformed data stored in the `ONLINE_BROWSING_HISTORY` table.

Extract life stage event information from social media posts

You can use social media to get valuable insights about customers. The Insurance sample uses the SMA Text Analytics model to extract information from social media posts.

Important: IBM Predictive Customer Intelligence does not retrieve social media data; this solution assumes that social media data is available.

The SMA Text Analytics model uses Text Analytics to read the social media data, and to extract life stage event information. Some examples of life stage events are new baby born, new job, new house, birthday, marriage, and so on. This information is used to help recommend appropriate insurance policies to customers.

Extract sentiment scores from customer complaints

Customer interactions with call center agents can be a source of valuable data to determine customer satisfaction levels. The Insurance sample uses the Sentiment Scoring model to extract the sentiment score from customer comments that are captured while recording customer complaints.

The input is customer complaint details. The example stream for extracting the sentiment score is named `Sentiment.str`. The Text Analytics model reads the customer complaint details, and extracts meaningful words and concepts from the information. Negative concepts are used to derive the sentiment score. The sentiment score reflects the count of negative words used by customers while making the complaints.

Example concepts are "bad", "not accessible", "slow", "wrong".

Insurance data model

The historic data that is used for predictive modeling in IBM Predictive Customer Intelligence is stored in an IBM DB2 database.

A Predictive Enterprise View (PEV) of the data is created which passes data to the real time scoring service. Database Views (DB Views) are also created to be used in IBM SPSS Modeler streams. The following table describes some of the data columns that are part of the Predictive Enterprise View and the Database Views.

Table 1. Key data columns in the Insurance sample

NAME	DESCRIPTION
AGE	The age of the policy holder.
CLTV	Customer Lifetime Value.
EDUCATION	The education level of the policy holder.
EMPLOYMENT_STATUS	The status of the person's employment.
GENDER	The person's sex or gender.
INCOME	The policy holder's annual income.
MARITAL_STATUS	The marital status of the person.
MAINTENANCE_COST	The cost of maintaining this policy.
MONTHS_SINCE_POLICY_INCEPTION	The number of months since the policy holder started the policy.
MONTHS_SINCE_LAST_CLAIM	The number of months since the policy holder filed the last claim.

Table 1. Key data columns in the Insurance sample (continued)

NAME	DESCRIPTION
NUMBER_OF_CLAIMS_DENIED	The number of claims that were denied.
NUMBER_OF_CLAIMS_FILED	The number of claims that are filed.
CLAIM_SETTLEMENT_DURATION	The time, in days between the date when the claim opened and the date when the claim was closed, and the customer satisfaction confirmed, based on the status of the claim.
NUMBER_OF_COMPLAINTS	The number of complaints the policy holder has submitted.
NO_OF_CLOSED_COMPLAINTS	The number of complaints that have been closed.
NUMBER_OF_OPEN_COMPLAINTS	The number of complaints that are open.
LATEST_NOTE_ATTITUDE	The noted attitude of the last communication.
AVG_NOTE_ATTITUDE	Average communication note attitude.
NUMBER_OF_POLICIES	The number of policies that the policy holder has.
POLICYHOLDER_ID	Any value without business meaning that uniquely distinguishes each occurrence of this entity.
POLICY_ID	Any value without business meaning that uniquely distinguishes each occurrence of this entity.
POLICY_TYPE	Indicates the policy type, for example, fixed term and flexible term.
VEHICLE_OWNERSHIP	Indicates whether the policy holder owns a vehicle or not.
VEHICLE_TYPE	The vehicle type.
VEHICLE_SIZE	The vehicle size.
HOME_OWNERSHIP_STATUS	The tenancy status of residence.
INSURANCE_LINES	The number of types of insurance products held by the policy holder.
INSURANCE_SCORE	The insurance score, based on the credit score, as well as other factors such as claim filing history.
LIFE_CUSTOMER	Indicates whether a customer owns a life insurance policy.
NON_LIFE_CUSTOMER	Indicates whether a person owns a non-life insurance policy.
NUMBER_OF_CHILDREN	The number of children of the policy holder.
NUMBER_OF_INSURED_CARS	The number of insured cars insured with the insurance company.
POLICYHOLDER_TENURE	The number of years the policy holder has been a customer with the insurance company.

Table 1. Key data columns in the Insurance sample (continued)

NAME	DESCRIPTION
TOTAL_PREMIUM	The total premium on all the policies paid by the policyholder to the insurance company.
RETIREMENT_PLAN	The name of the retirement plan owned by the policyholder
HOUSEHOLD_NUMBER_OF_CHILDREN	The number of children in household
HOUSEHOLD_NUMBER_OF_INSURED_CARS	The number of insured cars in the household that are insured with the insurance company.
NUMBER_OF_POLICIES_IN_HOUSEHOLD	The number of policies held by the household.
HOUSEHOLD_TENURE	The number of years that a household has been classified within its customer status.
HOUSEHOLD_PREMIUM	The total premium paid by the household to the insurance company.
HOUSEHOLD_DISPOSABLE_INCOME	The amount of money that the household can spend after having paid all the fixed expenses such as rent, mortgage repayment and so on.

Predictive models in the Energy and Utilities sample

A number of predictive models are provided in the Energy and Utilities sample, which you can install.

To install the samples, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

The following models form the basis of the predictive models in the Energy and Utilities sample:

Credit Rating Model

Predict the probability that customers will miss payments.

Demand Response Program Acceptance

Predict the probability that customers will need demand response programs.

Recommended Rate Plan

Recommend the most profitable Rate Plan.

Sentiment

Determine the sentiment score for customers based on social media information, such as Twitter.

Satisfaction

Measure customer satisfaction as determined by the net promoter score.

Predict credit rating

The Credit Rating model predicts the probability that customers will miss payments.

The example stream in the Energy and Utilities sample for predicting the credit rating for customers is named `Credit Rating Model.str`.

The Credit Rating model uses a TwoStep model to classify customers based on credit history, annual income, ownership, family members, and so on. Then a logistic regression algorithm is applied to read past payment history and customer classification to evaluate customers' ability to pay their bills.

Predict the need for demand response program assistance

The Demand Response Program Acceptance model predicts the probability that customers will need demand response programs.

The example stream in the Energy and Utilities sample is named `Demand Response Program Acceptance.str`.

A C5.0 algorithm reads the past history of demand response programs and customer demographic information and uses this information to predict whether customers will need demand response programs assistance in the future.

A demand response program encourages customers to reduce their energy use at times of peak demand to save money or gain incentives.

Recommend the right rate plan

The Recommended Rate Plan model in the Energy and Utilities sample recommends the most profitable rate plan.

The example stream in the Energy and Utilities sample is named `Recommended Rate Plan.str`. Average Energy and Utilities usage is considered by customer, and the bill amount is calculated for all available rate plans. The current rate plan bill totals are compared with all other rate plan bill totals, and from this comparison, the most profitable rate plan is selected.

Understand customer sentiment

The Sentiment model in the Energy and Utilities sample determines customer sentiment.

The example stream in the Energy and Utilities sample is named `Sentiment.str`.

An IBM SPSS Text Analytics model mines unstructured data to establish the customer sentiment. The data comes from social media sources such as Twitter. The data is categorized into areas such as outages, power line issues, customer service, and bill issues. A sentiment score is determined, based on the set of categories that is defined in the Text Analytics node.

Predictive models in the Banking sample

A number of predictive models are used in the Banking sample, which you can install.

To install the samples, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

The following models form the basis of the predictive models in the Banking sample:

Category affinity

Predicts what product or service the customer will be most interested in.

Churn Predicts whether customers are likely to renew their home insurance policy.

Credit card default

Predicts whether or not customers are likely to default on their credit card debt.

Customer segmentation

Segments customers into groups with similar demand characteristics, such as young educated and middle income, affluent and middle aged.

Sequence analysis

Predicts the kind of recommendation to make to each customer, based on what they have purchased. For example, a customer who obtains a mortgage is likely to want to purchase home insurance. A customer who has purchased travel insurance might want to activate a credit card for international use.

Determine category affinity in banking

You can determine customer affinity towards product lines by understanding the customer demographic information, purchase information, and browsing information.

The Category Affinity model uses customer transaction data as input (transaction date, merchant, category, and price) and predicts the customer's category affinity. It uses a logistic regression model. Logistic regression classifies records based on values of input fields.

The example stream for determining customer affinity towards product lines is named `Category Affinity.str`. The output is a predictive score that determines the probability that the customer will buy a product.

Predict churn in the banking industry

The churn model in the Banking sample predicts whether the customer is likely to renew their home insurance policy.

The example stream for predicting churn in the Banking sample is named `Churn.str`.

The churn model takes in several variables into account. For example, the customer's monthly premium, the number of months since the policy was started, the number of claims that are denied, the number of months for renewal, marital status, income, and age. It uses the CHAID algorithm.

Predict whether customers will default on credit card debt

The Credit Card Default model predicts whether or not customers are likely to default on credit card debt.

The example stream for predicting whether a customer will default on credit card debt is named `Credit Card Default.str`. The model uses customer information,

such as age, education, number of years employed, income, address, credit card debt, other debt, and whether the customer has defaulted in the past. The model uses a Bayesian Network algorithm.

Define customer segments in the banking industry

When you define customer segments, you profile the customers into groups that have similar demand characteristics.

The example stream for defining customer segments in the Banking sample is named `Customer_Segmentation.str`. The inputs into the model are customer information, such as age, education, years employed, address, income, and debt to income ratio. The model uses clustering from the TwoStep model to divide the customers into various segments, such as young educated and middle income, or affluent and middle aged. The TwoStep clustering method is used because it can handle mixed field types and is able to handle large data sets efficiently.

Sequence analysis

Sequence analysis predicts the kind of recommendation to make to each customer, based on what the customer has purchased. For example, a customer who obtains a mortgage is likely to want to purchase home insurance. A customer who has purchased travel insurance might want to activate a credit card for international use.

Sequence analysis uses a Sequence model, which discovers patterns in sequential or time-oriented data. The model detects frequent sequences and creates a generated model node that can be used to make predictions.

The example stream for sequence analysis is named `Sequence Analysis.str`.

Training predictive models

Predictive models must be trained to determine which data is useful and which data is not needed. When a model gives you accurate predictions, you are ready to use the predictive model for real time scoring.

You use a training data set to build the predictive model and a test set of data to validate the model that was created with the training set.

Models must be retrained periodically with new data sets to adjust for changing behavior patterns. For information about using IBM SPSS Modeler, see IBM SPSS Modeler Help (http://www-01.ibm.com/support/knowledgecenter/SS3RA7_16.0.0/com.ibm.spss.modeler.help/clementine/entities/clem_family_overview.htm?lang=en).

Scoring a model

To score a model means to apply it to some data in order to obtain a result or prediction that can be used as input to decisions.

Depending on the application, the scoring results can be written to a database table or flat file, or used as inputs to the segment, selection, and allocation rules that drive decisions in an application.

For more information, see IBM SPSS Collaboration and Deployment Services Deployment Manager User's Guide (<http://www.ibm.com/support/>

Creating business rules

You use IBM Analytical Decision Management to bring together your company's business rules, predictive models, and optimizations.

Using IBM Analytical Decision Management, the insights gained through predictive modeling can be translated to specific actions.

You can combine predictive models with rules to allocate offers in accordance with business goals. This is done using a combination of selection and allocation rules that are based on the output from predictive models.

The steps you take are:

Define possible actions

If a customer is not happy with a service, what should you do about it.

Allocate offers

Which types of customers are the best candidates for which offers.

Prioritize offers

Prioritization determines which offers a customer will receive.

For more information, see the IBM Analytical Decision Management Application User's Guide (http://www-01.ibm.com/support/knowledgecenter/SS6A3P_8.0.0/com.ibm.spss.dm.userguide.doc/configurableapps/dms_define_rules.htm).

Deployment

You can deploy the application to a testing environment or to a real-time production environment, such as a call center or a website. You can also deploy it to contribute to batch processing.

You can deploy a stream in the IBM SPSS Modeler repository. A deployed stream can be accessed by multiple users throughout the enterprise and can be automatically scored and refreshed. For example, a model can be automatically updated at regularly scheduled intervals as new data becomes available.

For more information, see IBM SPSS Collaboration and Deployment Services (http://www-01.ibm.com/support/knowledgecenter/SS69YH_6.0.0/com.spss.mgmt.content.help/model_management/_entities/whatsnew_overview_thick.html?cp=SS69YH_6.0.0%2F5).

Chapter 4. Integration with the IBM Enterprise Marketing Management suite

IBM Predictive Customer Intelligence integrates with the IBM Enterprise Marketing Management suite. The components that Predictive Customer Intelligence integrates with are IBM Campaign, IBM Interact and IBM Marketing Platform.

IBM Campaign

IBM Campaign is a web-based Enterprise Marketing Management (EMM) solution that enables users to design, run, and analyze direct marketing campaigns.

IBM Interact

With IBM Interact, users can retrieve personalized offers and customer profile information in real-time.

IBM Marketing Platform

The IBM Marketing Platform provides security, configuration, and dashboard features for IBM Enterprise Marketing Management products.

IBM Interact connectors

IBM Interact integrates with customer facing systems (touchpoints) such as web sites and call centers, and enables users to retrieve personalized offers and visitor profile information in real time.

IBM Predictive Customer Intelligence provides two connectors between IBM Interact and IBM SPSS Collaboration and Deployment Scoring Service:

- The External callout connector calls an SPSS model at runtime, and is contained within the expression of an advanced rule for a Marketer Score, overriding the score supplied by the IBM Enterprise Marketing Management campaign.
- The External learning connector extends IBM Interact's native learning module to monitor visitor actions and propose optimal offers. It prioritizes IBM Campaign offers based on an SPSS model's prediction of their final score. The connector passes specific configurable parameters as input to the SPSS Scoring Service.

The External Learning connector

The External Learning connector is used to override the marketer's score in IBM Interact with a customized score from IBM SPSS Collaboration and Deployment Services. The customized score is used to prioritize the offers for customers. When the customer facing touchpoint puts in an offer request to IBM Interact, it triggers the external learning connector. A touchpoint can be a call center web application, a mobile application, a batch command, and so on.

The following process is initiated:

1. The connector Java™ class calls the IBM SPSS Collaboration and Deployment Services scoring service that is based on the configuration setting of the current channel.
2. If, in the configuration file, the current channel is set to be monitored, the corresponding channel in the IBM SPSS Scoring Service is configured to update the channel offers' score number.

3. The External Learning connector uses IBM Interact runtime session profile data as SPSS Scoring Service input parameters.
4. The connector gets the scores from the SPSS scoring service and overwrites the Interact offers scores with the SPSS scores.
5. After the Interact offer's score is updated by the SPSS score, the offer list is sorted by the new final scoring number in descending order. The offer list is returned as the final offer list.




Figure 11. How the external learning connector works

The External Callout connector

By using the External Callout connector, you can override a single offer's score by calling an IBM SPSS Scoring Service. Also, using the External Callout connector, you can invoke the IBM SPSS score to present eligible offers.

To enable the External Callout connector, perform the following steps:

1. Register the External Callout connector in the IBM Enterprise Marketing Management (EMM) configuration.
2. Add the IBM SPSS Scoring Service information to the External Callout connector's configuration file.

Chapter 5. Integration with Next Best Action Optimizer

You can integrate IBM Next Best Action Optimizer (NBAOPT) with your IBM Predictive Customer Intelligence solution to further optimize your recommended actions.

NBAOPT segregates customers by lifetime value, and then provides recommendations to retain customers.

For information on installing the NBAOPT connector, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems* or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

Next Best Action Optimizer (NBAOPT)

Many companies across many industries (such as telecommunications, retail, and insurance) classify customers into different groups for marketing campaigns, targeted actions, and optimized use of resources. Classification helps Customer Relationship Management (CRM) departments to conduct promotional campaigns to retain customers, cross-sell, and up sell.

Traditional methods of classifying customers use standard customer and product segmentation strategies along with manually defined rules on an ad hoc basis. Often, the processes that govern these methods do not consider the following factors:

- Uncertainty in the behavior of the target entities
- Value and response to the prior actions
- The current state of the entity in relation to the organization
- Expected long-term rewards

NBAOPT addresses these problems by using a scalable data analytics and optimization engine. NBAOPT automatically generates an action allocation policy that is optimized for long term expected reward, subject to resource, legal and business constraints. The allocation policy can be used for targeted marketing, CRM, debt collection, retention, cross-sell, up-sell, and so on.

For more information, see *Next Best Action Optimizer (NBAOPT) Lifetime Value Maximizer & Collections Maximizer User Guide*.

How NBAOPT is used in a Predictive Customer Intelligence solution

NBAOPT works alongside IBM SPSS Modeler. When the NBAOPT connector is installed, NBAOPT connector items are added to the IBM SPSS Modeler menu so that you can run the NBAOPT Studio. In NBAOPT Studio, you set up the data for analysis. Then, you start the analysis, which performs the following actions:

- Segregates customers by lifetime value
- Provides recommended actions to retain customers, which are based on their lifetime value segment

This analytical process exports model nuggets in PMML format (Predictive Model Markup Language, an XML-based file format). Model nuggets can be used for scoring data for prediction, or analyzing data. From the PMML file that is created in NBAOPT Studio, you create an enhanced SPSS streaming SPSS Modeler that contains segmentation and recommendations. You can also generate visualizations of the results.

Open NBAOPT Studio from IBM SPSS Modeler

You can open IBM Next Best Action Optimizer (NBAOPT) Studio from IBM SPSS Modeler.

Before you begin

NBAOPT, and the NBAOPT connector must be installed. For information on installing the NBAOPT connector, see *IBM Predictive Customer Intelligence Installation Guide*.

Procedure

1. In IBM SPSS Modeler, click **Tools > NBA Optimizer**, and select **NBAOPT Studio**.
2. The first time that you start NBAOPT Studio, you are asked to specify the location of the NBAOPT Studio executable file. Click **Browse** to locate the file, and click **Select**.

What to do next

In NBAOPT Studio, you set up data for analysis and generate the PMML files that the NBAOPT connector creates the IBM SPSS streams from. For information on how to use NBAOPT Studio, see the *Next Best Action Optimizer (NBAOPT) Lifetime Value Maximizer & Collections Maximizer User Guide*, and the *Next Best Action Optimizer (NBAOPT) Lifetime Value Maximizer & Collections Maximizer Tutorial Guide*.

Generating streams from PMML files

IBM Next Best Action Optimizer (NBAOPT) model results are generated in PMML format, an XML format that can be used by a number of tools, including IBM SPSS Modeler. The NBAOPT engine generates SPSS-compliant PMML files for scoring.

Procedure

1. In IBM SPSS Modeler, click **Tools > NBA Optimizer**, and select **Generate Stream from the PMML Files**.
2. Browse for the PMML file location, for example: C:\NBAOPT\Configurations\demo_insurance\Modeling\Model1\SPSS. If multiple PMML files are present in the folder, multiple streams are generated.

The generated streams can have multiple branches. The number of branches depends on the number of action groups that are specified for a configuration. For example, if there are two actions, Channel and Product, two actions are specified during the scoring process.

3. Click **Generate Stream**.

In the example, two streams are generated and are named based on the PMML file names. The number of streams created is dependent on the number of PMML files.

These files are stored in `<model_folder>/SPSS`.

What to do next

After the stream is generated, and before you run the stream, configure the stream.

Example SPSS stream

This is an example of an IBM SPSS stream that is generated from PMML files




Figure 12. Example IBM SPSS stream generated from PMML files



Database source node

The database source node is used to import data from ODBC sources.



Model nuggets

A model nugget is a container for a model. It is used for scoring data to generate predictions or to allow further analysis of the model properties.

The nuggets are imported via PMML files that are generated from the NBAOPT engine.



Filler nodes

Filler nodes are used to replace field values and change storage.



Filter nodes

Filter nodes are used to filter (discard) fields, rename fields, and map fields from one source node to another.



Derive node

The derive node gives you the ability to create new fields from existing fields.



Merge node

The merge node merges multiple action group nugget branches.

For more information, see IBM SPSS Modeler Help (http://www-01.ibm.com/support/knowledgecenter/SS3RA7_16.0.0/com.ibm.spss.modeler.help/clementine/entities/clem_family_overview.htm?lang=en).

Configuring the IBM SPSS stream

When an IBM SPSS stream is generated from PMML files, you must configure the stream before you can deploy it.

Procedure

1. In IBM SPSS Modeler, open the stream to configure.
2. Double-click the database source node and add the **Data source** name and the **Table name** to the database source node (mandatory).
3. Verify the database settings by clicking **Preview**.
4. Double-click the filter nodes to exclude fields that are not to be included in the stream.
5. Double-click the merge node to set the merge key (mandatory if there is a merge node). In the **Merge** tab, move fields into the **Keys for merge** list.
6. Set the filler node to replace database blank values (optional). For example, you can replace blank values with an empty string, such as `to_string(@FIELD)`.
7. Add extra nodes to the stream if required.

For example, if a field has a different data type in the database table, and the NBAOPT Studio Mining Type configuration, you must add a derive node to convert the character type to numeric.

Check that the derive field name matches the model nugget field name, if it does not, change the derive field name to match.

8. Run the stream to verify the stream output table. In IBM SPSS Modeler, from the **Tools** menu, click **Run**.

For more information, see IBM SPSS Modeler Help (http://www-01.ibm.com/support/knowledgecenter/SS3RA7_16.0.0/com.ibm.spss.modeler.help/clementine/buildingstreams_container.htm?lang=en).

What to do next

Add the configured stream to the IBM Collaboration and Deployment Services repository file, and configure the scoring service. For more information, see IBM SPSS Collaboration and Deployment Services User Guide (http://www-01.ibm.com/support/knowledgecenter/SS69YH_6.0.0/com.spss.mgmt.content.help/model_management/thick/scoring_configuration_overview.html?lang=en).

After the scoring service is configured, the IBM Collaboration and Deployment Services can access the scoring service, and other applications can use the scoring service to get the next best action based on the input fields values.

Chapter 6. Reports and portal pages

You can install a number of sample reports and portal pages, and modify them for your own purposes. You can also install the IBM Predictive Customer Intelligence Usage Report. The Usage Report monitors the number of offers or recommendations that are provided.

Sample reports and portal pages

A number of Predictive Customer Intelligence reports and portal pages can be installed with the samples.

The following industry samples include reports: Telecommunications, Retail, Insurance, and Energy and Utilities. The Banking sample does not include reports.

For information about installing the sample reports, see the *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*.

You can customize the reports and portal pages using IBM Cognos Report Studio. Cognos Report Studio is a report design and authoring tool. Report authors can use Report Studio to create, edit, and distribute a wide range of professional reports. For more information see IBM Cognos Report Studio User Guide (http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.1/com.ibm.swg.ba.cognos.ug_cr_rptstd.10.2.1.doc/c_rs_introduction.html).

The metadata that the report displays comes from the package that is created in IBM Cognos Framework Manager. The example Framework Manager project folder contains the compiled project file (.cpf). When you open the .cpf file, Framework Manager displays the modeled relationships of the data and the package definitions, which are made available to the reporting studios when published. You can modify the metadata for the report by using Framework Manager. For more information, see IBM Cognos Framework Manager User Guide (http://www-01.ibm.com/support/knowledgecenter/SSEP7J_10.2.1/com.ibm.swg.ba.cognos.ug_fm.10.2.1.doc/c_ug_fm_introduction.html%23ug_fm_Introduction).

Reports available in the Telecommunications, and Energy and Utilities samples

The following Predictive Customer Intelligence reports are available from the home page for the Telecommunications, and Energy and Utilities samples.

BillingReports

The **Billing** report displays the customer billing history, calculated bill amount, and potential savings.

CaseDetailReports

The **CaseReport** shows the open and closed case history of the customer: a bar chart that plots the case status per month, a list with detail case information and a list that shows social media posts.

NetworkReports

The **StoresMapReport** shows the location of nearest stores for a specified location, in a map report.

ProfileReports

The **Dials** report shows the key analytics key performance indicators (KPIs) for a customer.

For the Telecommunications sample only, the **SocialNetworkChart** report shows the social network chart for a customer ID.

UsageReports

The **Usage** report shows the voice and data usage of the customer.

Portal pages available in the Retail sample

The following Predictive Customer Intelligence portal pages are available from the home page for the Retail sample.

Customer Online Activity Analysis

Shows detailed analysis of Customer's online activities with the following charts:

- Purchase Activity Analysis (Traditional) shows a distribution of purchase activities across different product lines.
- All Online Activity Analysis shows a relative comparison of all online activities for all product lines
- Relative Activity Conversion Analysis shows a stacked chart to analyze the conversions between different online activities such as Page Views, Products Browsed, Products Carted, Products Abandoned, Products Purchased, for all product lines.
- Analyzing Online Activities by Customer Segments allows you to analyze the online activities for different customer segments and combination of customer segments.
- Analyzing Online Activities for Demographic Cluster shows a detailed analysis of online activities for a demographic cluster. For example, for a selected product line which cluster has highest products browsed activity.

Customer Response Analysis

Contains two charts that are based on the data available from customer's response to product satisfaction and various offers that were made through the IBM Analytical Decision Management application.

Shows a dial chart that shows how the customers rated a particular product line or a product. It also contains a column chart that shows the relative comparison of the probability of accepting different offers.

Market Basket Analysis

Shows the products that are likely to be bought together. This report can also be filtered for individual products.

Understanding the Customer

Shows a distribution of customer's online activities across different product lines, and customer segmentation. Customer segments are shown by using customer's demographic data, online activity data, and their purchasing behavior.

Reports available in the Insurance sample

The following Predictive Customer Intelligence reports are available from the home page for the Insurance sample.

CaseReportInsurance

Shows the cases by month for the selected client.

Churn and Sentiment Scores

The Dial report shows churn propensity and sentiment score.

Churn Propensity

Shows the churn propensity ratio (the willingness of the client to leave the provider) for the selected client.

ClaimsReport

The list report shows customer claims information.

Complaints

The bar chart and list report shows customer complaints information.

PolicyDetails

The list report shows customer policy information.

SMA Insight

Shows the social media activity for the selected client.

Social and Web Analytics

The pie chart and list report shows customer social media post and insurance product browsing summary information.

Viewing the sample reports

If the sample reports are installed for Telecommunications (Telco), Retail, Insurance, and Energy and Utilities, you can view and test them.

Procedure

1. In your browser, type the URL for IBM Cognos Connection, which is supplied by your administrator, and then press **Enter**. The URL looks something like this: `http://servername/IBMCognos`.
2. In the **IBM Cognos Welcome** page, click **IBM Cognos content**.
3. In the **Public Folders** tab, click: **IBM PCI < Sample_name >** to list the available reports or portal.
For Retail, click **IBM PCI Retail Portal Pages** to list the reports.
Insurance reports are listed within an **Insurance** subfolder.
4. You can use the following test customer IDs for samples where a Customer ID is required.
 - For the Telecommunications, and Energy and Utilities sample, type 21.
 - For the Insurance sample, use the following inputs for the specified reports:
 - CaseReportInsurance, Churn and Sentiment Scores, Churn Propensity, ClaimsReport, and PolicyDetails: 1001920.
 - SMA Insight: AaronBBailey@trashymail.com
 - Social and Web Analytics: 1006598

Modifying the Predictive Customer Intelligence data model

IBM Predictive Customer Intelligence uses IBM Cognos Framework Manager to model the metadata for reports.

Before you begin

The computer where IBM Cognos Framework Manager is installed must have access to the folder that contains the IBM Predictive Customer Intelligence project files. For more information, see the *IBM Predictive Customer Intelligence Installation Guide*.

About this task

IBM Cognos Framework Manager is a metadata modeling tool that drives query generation for IBM Cognos software. A model is a collection of metadata that includes physical information and business information for one or more data sources.

For information on modifying or creating Framework Manager models, see the *IBM Cognos Framework Manager User Guide*. This is available from IBM Knowledge Center (http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.1/com.ibm.swg.ba.cognos.cbi.doc/welcome.html).

Procedure

1. Launch IBM Cognos Framework Manager.
2. Click **Open a Project** and browse to the folder containing the .cpf file. For example, IBM PCI Telco.cpf.

Predictive Customer Intelligence Usage Report

The IBM Predictive Customer Intelligence Usage Report monitors the number of offers or recommendations that are provided. The report also records the response rates that are broken down by type of response.

The number of offers that are presented, accepted, and rejected are shown by channel (for example, website, mobile application) and by month. The numbers of offers that are purchased are shown by month.

If you use IBM Enterprise Marketing Management (EMM) as the recommendation generator, the data comes from the system tables that are used for logging offers. If you use IBM Analytical Decision Management as the recommendation generator, only the number of offers that are presented is found in the log tables. In this case, if you want to capture acceptance and rejection, you can build custom extensions to your call center or web application.

The Usage Report is installed as part of the artifacts installation. For more information, see the *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

For information about populating the Predictive Customer Intelligence database for the Usage Report, see Appendix B, "Configuration of the Usage Report," on page 73.

You can customize the report using IBM Cognos Report Studio. Cognos Report Studio is a report design and authoring tool. Report authors can use Report Studio to create, edit, and distribute a wide range of professional reports. For more information about how to use Report Studio, see the *IBM Cognos Report Studio User Guide*. You can obtain this user guide from IBM Knowledge Center,

(http://www.ibm.com/support/knowledgecenter/SSEP7J_10.2.1/com.ibm.swg.ba.cognos.ug_cr_rptstd.10.2.1.doc/c_rs_introduction.html).

The metadata that the report displays comes from the package that is created in IBM Cognos Framework Manager. The example Framework Manager project folder contains the compiled project file (.cpf). When you open the .cpf file, Framework Manager displays the modeled relationships of the data and the package definitions, which are made available to the reporting studios when published. You can modify the metadata for the report by using Framework Manager. For more information, see “Modifying the Predictive Customer Intelligence data model” on page 61.

Running the Predictive Customer Intelligence Usage Report

You can run the IBM Predictive Customer Intelligence Usage Report from within IBM Cognos Connection. You can also filter the data that is displayed in the report. By default, the report is not filtered.


Procedure

1. Open IBM Cognos Connection, and navigate to the location of the report. By default, this is **Public Folders > IBM Predictive Customer Intelligence > Reports**.
2. Click the IBM Predictive Customer Intelligence Usage Report link. By default, the report is rendered in IBM Cognos Viewer in HTML format.
3. You can filter the data that is displayed in the **Recommendations by Channel** crosstab by making selections in the prompt controls and then clicking **Refresh** to update the report. You can filter by Date, Response type, or by Channel. To hide the prompt controls click **Filters**. The prompt values selected are displayed when the controls are hidden, to ensure that the report context is maintained.
4. You can print the report in PDF format. On the Cognos Viewer toolbar, click the **Output** icon and select **View in PDF format** (the icon displays the current output format, such as HTML).

Accessibility for the Usage Report

The Usage Report has been designed to be accessible for those users who require the use of such technologies.

To enable accessibility options, do the following steps:

1. In IBM Cognos Connection, click **Run with options**  adjacent to the report.
2. Select **Enable accessibility support**, and click **Run**.

The report is rendered with the accessibility support enabled. This results in adjustments to the report layout to accommodate this support.

The Framework Manager model for the Usage Report

The IBM Cognos Framework Manager model contains the metadata for the Usage Report.

A view of the database layer in the IBM Cognos Framework Manager Content Explorer shows the relationships between the database tables.

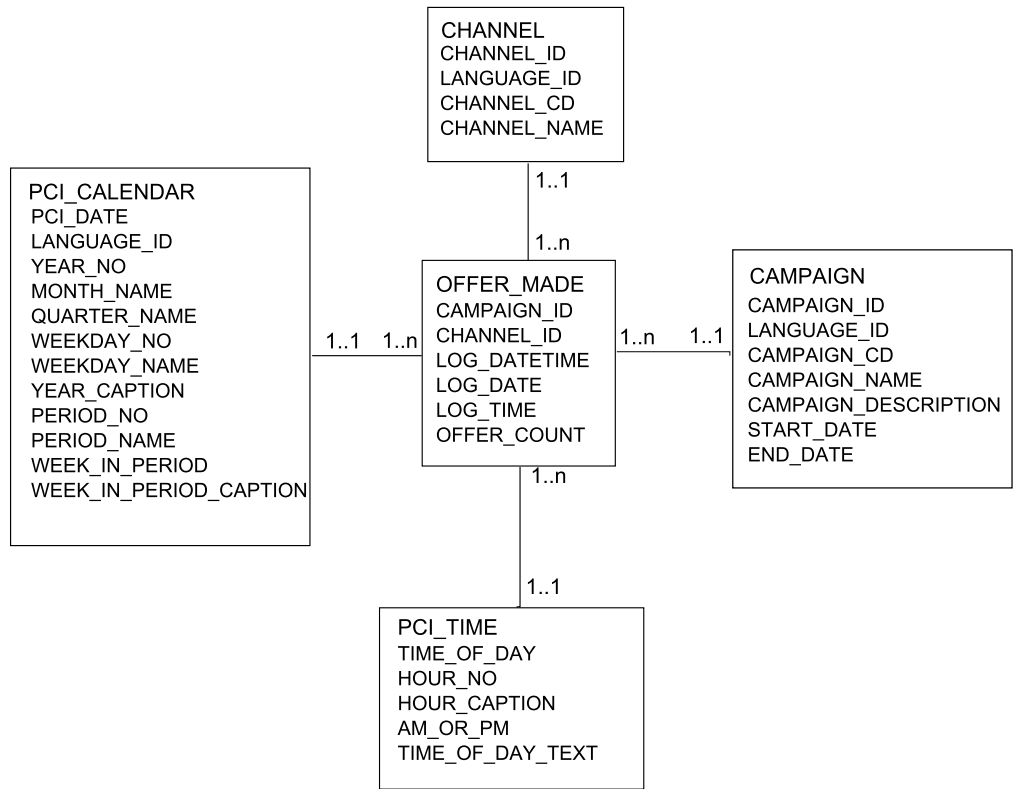


Figure 14. A view of the Offer Made tables

A view of Offer Responses that shows the relationships between the tables responsible for the Offer Responses that are shown in the Usage Report is shown in the following figure.

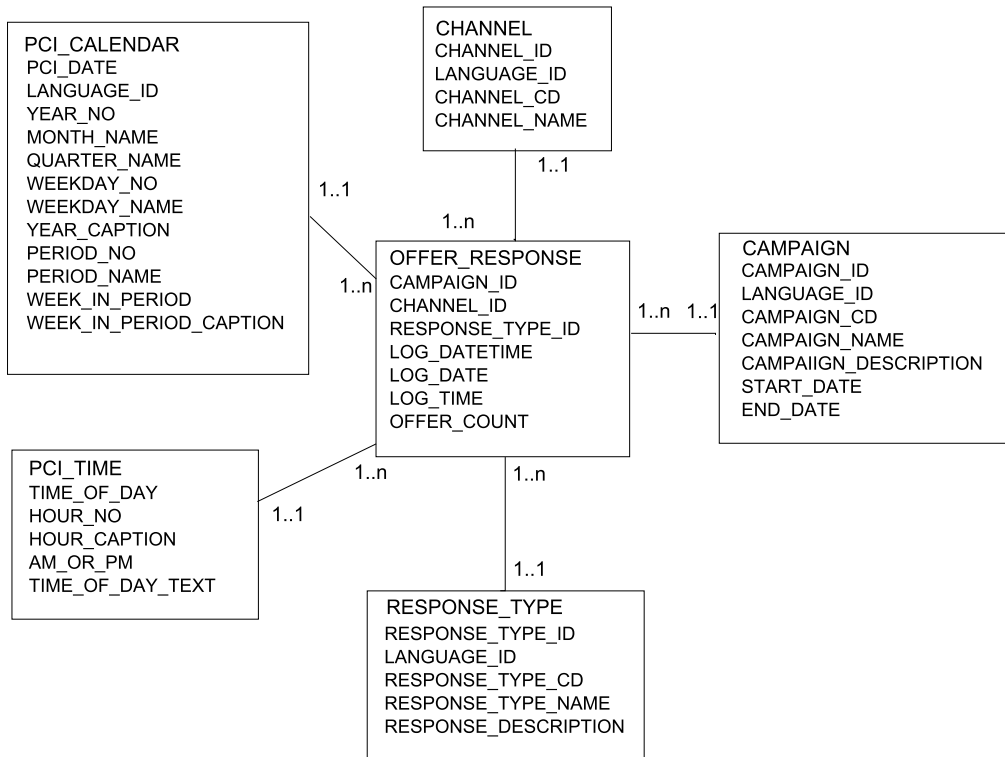


Figure 15. A view of the Offer Response tables

Appendix A. IBM Predictive Customer Intelligence samples

Samples, including model files, reports, and other content, is supplied to help you to configure IBM Predictive Customer Intelligence for the needs of your business.

Telecommunications sample

You can install a sample that demonstrate how IBM Predictive Customer Intelligence can be used to retain customers in the Telecommunications (Telco) industry.

To see a case study that illustrates the Telecommunications sample, see “Telecommunications case study” on page 3.

The following table provides a summary of the Telecommunications sample artifacts that are provided with IBM Predictive Customer Intelligence.

Table 2. Telecommunications sample content

Type of sample content	File names	Content
Business Intelligence Reports	IBM_PCI_Telco_Cognos_Content.zip IBM_PCI_Telco_Model.zip IBM_PCI_Telco_Report_Images.zip	The following reports are included: <ul style="list-style-type: none">• Billing reports• Case Detail reports• Networks reports• Profile reports:<ul style="list-style-type: none">– Dials report– Social Network Chart report• Usage reports The reports are described in “Sample reports and portal pages” on page 59.
IBM Analytical Decision Management applications and templates	IBM_PCI_Telco_App.zip TelcoCallCenter.xml	For more information, see “Telecommunications models in Analytical Decision Management” on page 30
Predictive Models	IBM_PCI_Telco.pes	The following predictive models are included: <ul style="list-style-type: none">• Churn model• Customer satisfaction model• Association model• Response propensity model The Predictive Models are described in “Predictive models in the Telecommunications sample” on page 26.

Table 2. Telecommunications sample content (continued)

Type of sample content	File names	Content
DB2 database	IBM_PCI_Telco_Data.zip	

To install the samples, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

Retail sample

You can install a sample that demonstrate how IBM Predictive Customer Intelligence can be used to increase revenue through targeted recommendations to customers.

To see a case study that illustrates the Retail sample, see “Retail case study” on page 11.

The following table is a summary of the Retail sample artifacts that are provided with IBM Predictive Customer Intelligence.

Table 3. Retail sample content

Type of sample content	File names	Content
Business Intelligence Reports	IBM_PCI_Retail_Cognos_Content.zip IBM_PCI_Retail_Model.zip	The following reports are included: <ul style="list-style-type: none"> • Understanding the Customer • Market Basket Analysis • Customer Online Activity Analysis <ul style="list-style-type: none"> – Purchase Activity Analysis – All Online Activity Analysis – Relative Activity Conversion Analysis – Analyzing Online Activities by Customer Segments – Analyzing Online Activities by Demographic Cluster • Customer Response Analysis <p>The reports are described in “Sample reports and portal pages” on page 59.</p>
IBM Analytical Decision Management applications and templates	IBM_PCI_Retail_App.zip Retail.xml	The Analytical Decision Management applications and templates are described in “Use Retail case study models in IBM Analytical Decision Management” on page 39.

Table 3. Retail sample content (continued)

Type of sample content	File names	Content
Predictive Models	IBM_PCI_Retail.pes	<p>The following predictive models are included:</p> <ul style="list-style-type: none"> • Customer segmentation model • Market basket analysis model • Customer affinity model • Response log analysis model • Price sensitivity model • Inventory-based suggestion model <p>The Predictive Models are described in “Predictive models in the Retail sample” on page 31.</p>
DB2 database	IBM_PCI_Retail_Data.zip	

To install the samples, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

Insurance sample

You can install a sample that demonstrate how IBM Predictive Customer Intelligence can be used to deliver the most appropriate action to each insurance customer during each interaction by analyzing customer data.

The following table provides a summary of the Insurance sample artifacts that are provided with IBM Predictive Customer Intelligence.

Table 4. Insurance sample content

Type of sample content	File names	Content
Business Intelligence Reports	<p>IBM PCI Insurance Cognos Content.zip</p> <p>IBM PCI Insurance Model.zip</p>	<p>The following reports are included:</p> <ul style="list-style-type: none"> • CaseReportInsurance • Churn and Sentiment Scores • Churn Propensity • ClaimsReport • Complaints • PolicyDetails • SMA Insight • Social and Web Analytics <p>The reports are described in “Sample reports and portal pages” on page 59.</p>

Table 4. Insurance sample content (continued)

Type of sample content	File names	Content
IBM Analytical Decision Management applications and templates	IBM_PCI_Insurance_App.zip Insurance.xml	
Predictive Models	IBM_PCI_Insurance.pes	<p>The following predictive models are included:</p> <ul style="list-style-type: none"> • Segmentation model • Churn prediction model • Customer lifetime value model (CLTV) • Campaign response model • Auto churn model • Life stage segment model • Buying propensity model • Data processing stream • Insurance policy recommendation model • Social Media Analytics (SMA) model • Sentiment scoring model <p>The Predictive Models are described in “Predictive models in the Insurance sample” on page 40.</p>
DB2 database	IBM_PCI_Insurance_Data.zip	The insurance data model is described in “Insurance data model” on page 46.

To install the samples, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

Banking sample

You can install a sample that demonstrate how IBM Predictive Customer Intelligence can be used in the Banking industry.

The following table provides a summary of the Banking sample artifacts that are provided with IBM Predictive Customer Intelligence.

Table 5. Banking sample content

Type of sample content	File names	Content
Predictive Models	IBM_PCI_Banking.pes	<p>The following predictive models are included:</p> <ul style="list-style-type: none"> • Category affinity • Churn • Credit card default • Customer segmentation • Sequence analysis <p>The Predictive Models are described in “Predictive models in the Banking sample” on page 49.</p>
DB2 database	IBM_PCI_Banking_Data.zip	

To install the samples, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

Energy and Utilities sample

You can install a sample that demonstrate how IBM Predictive Customer Intelligence can be used by Energy and Utilities organizations to retain customers and increase customer satisfaction by proactively resolving customer issues.

To see a case study that describes the Energy and Utilities sample, see “Energy and Utilities case study” on page 21.

The following table is a summary of the Energy and Utilities sample artifacts that are provided with IBM Predictive Customer Intelligence.

Table 6. Energy and Utilities sample content

Type of sample content	File names	Content
Business Intelligence Reports	<p>IBM PCI Energy and Utilities Cognos Content.zip</p> <p>IBM PCI Energy and Utilities Model.zip</p>	<p>The following reports are included:</p> <ul style="list-style-type: none"> • Billing reports • Case Detail reports • Networks reports • Profile reports: <ul style="list-style-type: none"> – Dials report (Key performance indicators for a customer) • Usage reports <p>The reports are described in “Sample reports and portal pages” on page 59.</p>
IBM Analytical Decision Management applications and templates	<p>IBM_PCI_Energy_and_Utilities_App.zip</p> <p>EnergyandUtilities.xml</p>	

Table 6. Energy and Utilities sample content (continued)

Type of sample content	File names	Content
Predictive Models	IBM_PCI_EU.pes	The following predictive models are included: <ul style="list-style-type: none"> • Credit Rating model • Demand Response Program Acceptance • Recommended Rate Plan • Sentiment • Satisfaction
DB2 database	IBM_PCI_EU_Data.zip	

To install the samples, see *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*.

Appendix B. Configuration of the Usage Report

The IBM Predictive Customer Intelligence Usage Report monitors the number of offers or recommendations that are provided.

If you use IBM Enterprise Marketing Management as the recommendation generator, the data comes from the system tables that are used for logging offers. If you use IBM Analytical Decision Management as the recommendation generator, only the number of offers that are presented is found in the log tables. In this case, if you want to capture acceptance and rejection, build custom extensions to your call center or web application.

To configure the Usage Report, you must first configure the logging of events, and then you must populate the Predictive Customer Intelligence database. The steps to do this differ depending on whether you are using Enterprise Marketing Management or Analytical Decision Management as the recommendation generator.

The Usage Report is installed as part of the Artifacts installation. For more information, see the *IBM Predictive Customer Intelligence Installation Guide for Microsoft Windows Operating Systems*, or *IBM Predictive Customer Intelligence Installation Guide for Linux Operating Systems*, available from IBM Knowledge Center (http://www.ibm.com/support/knowledgecenter/SSCJHT_1.0.0).

Predictive Customer Intelligence database tables

The IBM Predictive Customer Intelligence custom database contains the following 11 tables and their attributes.

CAMPAIGN

The CAMPAIGN master data table contains the campaigns that an offer belongs to.

Table 7. CAMPAIGN

Column	Data Type
CAMPAIGN_ID	INTEGER(4)
LANGUAGE_ID	INTEGER(4)
CAMPAIGN_CD	VARGRAPHIC(50)
CAMPAIGN_NAME	VARGRAPHIC(200)
CAMPAIGN_DESCRIPTION	VARGRAPHIC(500)
START_DATE	DATE(4)
END_DATE	DATE(4)

CHANNEL

The CHANNEL master data table contains the communication channels that interact with customers.

Table 8. CHANNEL

Column	Data Type
CHANNEL_ID	INTEGER(4)
LANGUAGE_ID	INTEGER(4)
CHANNEL_CD	VARGRAPHIC(50)
CHANNEL_NAME	VARGRAPHIC(200)

KEY_LOOKUP

The KEY_LOOKUP master data table contains the foreign keys.

Table 9. KEY_LOOKUP

Column	Data Type
KEY_LOOKUP_ID	BIGINT(8)
TABLE_NAME	VARGRAPHIC(50)
KEY_LOOKUP_CD	VARGRAPHIC(50)

PCI_CALENDAR

The PCI_CALENDAR master data table contains the calendar.

Table 10. PCI_CALENDAR

Column	Data Type
PCI_DATE	DATE(4)
LANGUAGE_ID	INTEGER(4)
YEAR_NO	INTEGER(4)
MONTH_NO	INTEGER(4)
QUARTER_NO	INTEGER(4)
MONTH_NAME	VARGRAPHIC(20)
QUARTER_NAME	VARGRAPHIC(20)
WEEKDAY_NO	INTEGER(4)
WEEKDAY_NAME	VARGRAPHIC(10)
YEAR_CAPTION	VARGRAPHIC(10)
PERIOD_NO	INTEGER(4)
PERIOD_NAME	VARGRAPHIC(25)
WEEK_IN_PERIOD	INTEGER(4)
WEEK_IN_PERIOD_CAPTION	VARGRAPHIC(25)

PCI_LANGUAGE

The PCI_LANGUAGE master data table contains the language codes used for globalization.

Table 11. PCI_LANGUAGE

Column	Data Type
LANGUAGE_ID	INTEGER(4)
LANGUAGE_CD	VARGRAPHIC(50)
LANGUAGE_NAME	VARGRAPHIC(50)

PCI_TIME

This master data table contains the time, down to the second.

Table 12. PCI_TIME

Column	Data Type
TIME_OF_DAY	TIME(3)
HOUR_NO	INTEGER(4)
HOUR_CAPTION	VARCHAR(5)
AM_OR_PM	VARGRAPHIC(25)
TIME_OF_DAY_TEXT	VARCHAR(50)

OFFER_MADE

The OFFER_MADE fact table records the number of offers made by date and time.

Table 13. OFFER_MADE

Column	Data Type
CAMPAIGN_ID	INTEGER(4)
CHANNEL_ID	INTEGER(4)
LOG_DATETIME	TIMESTAMP(10)
LOG_DATE	DATE(4)
LOG_TIME	TIME(3)
OFFER_COUNT	INTEGER(4)

OFFER_RESPONSE

The OFFER_RESPONSE fact table records the number of responses received by type, date and time.

Table 14. OFFER_RESPONSE

Column	Data Type
CAMPAIGN_ID	INTEGER(4)
CHANNEL_ID	INTEGER(4)
RESPONSE_TYPE_ID	INTEGER(4)
LOG_DATETIME	TIMESTAMP(10)
LOG_DATE	DATE(4)
LOG_TIME	TIME(3)
OFFER_COUNT	INTEGER(4)

OFFER_TARGET_MONTH

The OFFER_TARGET_MONTH fact table contains a record of how many recommendations were purchased by month in a given year. It is populated during installation if required.

Usually each row in OFFER_TARGET_MONTH has one twelfth of the recommendations in OFFER_TARGET_YEAR for the same year, but that value can be overridden.

Table 15. OFFER_TARGET_MONTH

Column	Data Type
PURCHASE_YEAR	INTEGER(4)
PURCHASE_MONTH	INTEGER(4)
RECOMMENDATION_COUNT	INTEGER(4)

OFFER_TARGET_YEAR

The OFFER_TARGET_YEAR fact table contains a record of how many recommendations were purchased by year. It is populated during installation if required.

Table 16. OFFER_TARGET_YEAR

Column	Data Type
PURCHASE_YEAR	INTEGER(4)
RECOMMENDATION_COUNT	INTEGER(4)

RESPONSE_TYPE

The RESPONSE_TYPE master data table contains the range of response types to an offer.

Table 17. RESPONSE_TYPE

Column	Data Type
RESPONSE_TYPE_ID	INTEGER(4)
LANGUAGE_ID	INTEGER(4)
RESPONSE_TYPE_CD	VARGRAPHIC(50)
RESPONSE_TYPE_NAME	VARGRAPHIC(200)

Configuring logging in IBM Enterprise Marketing Management

If you use IBM Enterprise Marketing Management as the recommendation generator, and you are using the IBM Predictive Customer Intelligence Usage Report, you must configure the logging of events for different categories.

About this task

Communication channels are configurable in IBM Enterprise Marketing Management. Part of the configuration includes setting up logging of events for different categories. The default category of **Get Offer** must be logged for

acceptance and rejection. If there are user-defined categories for acceptance and rejection, they must also be set to log for acceptance or rejection.

Procedure

1. Log in as administrator to the **IBM Campaign Manager** console.
2. Select **Campaign**, and then **Interactive Channels**.
3. Select and edit each interactive channel:
 - a. Click the **Events** tab.
 - b. Select the event **Get Offer**, or any other user-defined event for accepting or rejecting an offer.
 - c. Select **Log Offer Made**, **Log Acceptance**, and **Log Rejection**.

Populate the Predictive Customer Intelligence database from Enterprise Marketing Management

If you use IBM Enterprise Marketing Management as the recommendation generator, the data for the IBM Predictive Customer Intelligence Usage Report comes from the system tables that are used for logging offers.

The following tables show how the data must be mapped between the IBM Predictive Customer Intelligence database, and the IBM Enterprise Marketing Management database.

Table 18. Map the UA_CAMPAIGN table to the CAMPAIGN table

Predictive Customer Intelligence Column: CAMPAIGN	Enterprise Marketing Management Column: UA_CAMPAIGN	Transformations
CAMPAIGN_ID	<Assigned by system>	
CAMPAIGN_CD	CAMPAIGNID	Convert number to vargraphic
CAMPAIGN_NAME	NAME	
CAMPAIGN_DESCRIPTION	DESCRIPTION	
START_DATE	STARTDATE	
END_DATE	ENDDATE	

Table 19. Map the UACI_INTCHANNEL table to the CHANNEL table

Predictive Customer Intelligence Column: CHANNEL	Enterprise Marketing Management Column: UACI_INTCHANNEL	Transformations
CHANNEL_ID	ICID	
CHANNEL_CD	ICID	Convert number to vargraphic
CHANNEL_NAME	NAME	

Table 20. Map the UACI_EVENTACTIVITY table to the OFFER_MADE table

Predictive Customer Intelligence Column: OFFER_MADE	Enterprise Marketing Management Column: UACI_EVENTACTIVITY table	Transformations
CHANNEL_ID	ICID	

Table 20. Map the UACI_EVENTACTIVITY table to the OFFER_MADE table (continued)

Predictive Customer Intelligence Column: OFFER_MADE	Enterprise Marketing Management Column: UACI_EVENTACTIVITY table	Transformations
CAMPAIGN_ID	N/A	
OFFER_COUNT	OCCURRENCES	Sum
LOG_DATETIME	DATEID + TIMEID	Concatenate and convert to datetime

Filters:

- EVENTNAME = 'Get Offer'
- Group by ICID, DATEID, TIMEID

Table 21. Map the UACI_EVENTACTIVITY table to the OFFER_RESPONSE table

Predictive Customer Intelligence Column: OFFER_RESPONSE table	Enterprise Marketing Management Column: UACI_EVENTACTIVITY	Transformations
CHANNEL_ID	ICID	
CAMPAIGN_ID	N/A	
RESPONSE_TYPE_ID	EVENTID	
OFFER_COUNT	OCCURRENCES	Sum
LOG_DATETIME	DATEID + TIMEID	Concatenate and convert to datetime

Filters:

- CATEGORYNAME = 'Response'
- Group by ICID, DATEID, TIMEID

Table 22. Map the UACI_EVENTACTIVITY table to the RESPONSE_TYPE table

Predictive Customer Intelligence Column: RESPONSE_TYPE	IBM Enterprise Marketing Management Column: UACI_EVENTACTIVITY	Transformations
RESPONSE_TYPE_ID	EVENTID	
RESPONSE_TYPE_CODE	EVENTID	Convert number to vargraphic
RESPONSE_TYPE_NAME	EVENTNAME	

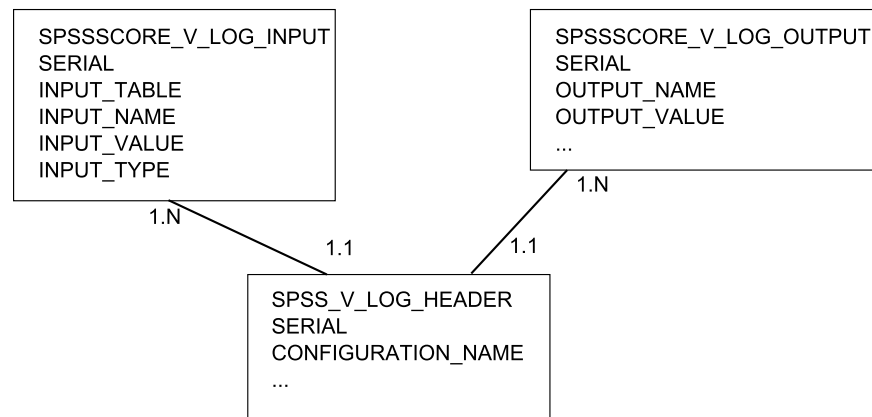
Filters:

- distinct EVENTID
- CATEGORYNAME= 'Response'

Configure logging in IBM SPSS Collaboration and Deployment

If you use IBM Analytical Decision Management as the recommendation generator, and you are using the IBM Predictive Customer Intelligence Usage Report, you cannot get the campaign, channel, customer, offer, or response from the IBM SPSS database tables. This information must be obtained from another application, and then used as inputs to the decision model to be available for logging

The following views are available in IBM SPSS Collaboration and Deployment Services when logging is configured:



SPSSCORE_V_LOG_HEADER contains the scoring models that are configured for real-time decisions.

SPSSCORE_V_LOG_INPUT contains the attributes that are used in assessing the scenario and rendering a recommendation.

SPSSCORE_V_LOG_OUTPUT contains the attributes that are returned from the decision model, which can include some of the inputs and the recommendation.

Configure logging in SPSS

You can configure logging in SPSS down to the attribute level. Consider the following points:

- The channel must be an input field to the model and must be set up for logging.
- The campaign must be an input field to the model and must be set up for logging.
- Any other dimensions that are required by the dashboard, such as campaign, must be both inputs and logged outputs of the model.

Configure a model for scoring by using IBM SPSS Deployment Manager. During scoring configuration, you can select any input or output field for logging. The customer data determines what is available for logging.

For example, for the Telecommunications sample, select the following fields for logging:

- CALL_CENTER_RESPONSE
- DIRECT_MAIL_RESPONSE
- EMAIL_RESPONSE
- SMS_RESPONSE
- CURRENT_OFFER

Select the following **Model Outputs**:

- Campaign
- Offer
- Output-PredictedProfit
- Output-MaxOffersNum
- Output-MinProfit

- Output-ProbToRespond
- Output-Revenue
- Output-Cost

For more information, see IBM SPSS Deployment Manager User's Guide (http://www-01.ibm.com/support/knowledgecenter/SS69YH_6.0.0/com.spss.mgmt.content.help/model_management/thick/idh_dlg_scoring_configuration_logging.html).

Populate the Predictive Customer Intelligence database from IBM SPSS

If you use IBM Analytical Decision Management as the recommendation generator for the IBM Predictive Customer Intelligence Usage Report, only the number of offers that are presented is found in the log tables.

In IBM SPSS, there are no dedicated system tables for channel, response type, and offers. Custom database tables must be used for the channel, response type, and offers.

The following tables show how the data must be mapped between the IBM Predictive Customer Intelligence database, and IBM SPSS database tables.

Table 23. CAMPAIGN table for Predictive Customer Intelligence mapped to the custom database tables from IBM SPSS

Predictive Customer Intelligence Column	SPSS Column	Filters
CAMPAIGN_ID	Sequentially generated number	These are distinct rows because Campaigns do not repeat.
CAMPAIGN_CD	I.INPUT_VALUE	H.CONFIGURATION_NAME = 'name of model'; I.INPUT_NAME = 'campaign cd';
CAMPAIGN_NAME	I.INPUT_VALUE	H.CONFIGURATION_NAME = 'name of model'; I.INPUT_NAME = 'campaign name';

Replace the filters in quotation marks with the names for the database attributes.

SPSS view:

```
SPSSSCORE_V_LOG_HEADER AS H
Join
SPSSSCORE_V_LOG_INPUT AS I on H.SERIAL = I.SERIAL
```

Table 24. CHANNEL table for Predictive Customer Intelligence mapped to the custom database tables from IBM SPSS

Predictive Customer Intelligence Column	SPSS Column	Filters
CHANNEL_ID	Sequentially generated number	Distinct rows so channels do not repeat

Table 24. CHANNEL table for Predictive Customer Intelligence mapped to the custom database tables from IBM SPSS (continued)

Predictive Customer Intelligence Column	SPSS Column	Filters
CHANNEL_CD	I.INPUT_VALUE	H.CONFIGURATION_NAME = 'name of model'; I.INPUT_NAME = 'channel cd';
CHANNEL_NAME	I.INPUT_VALUE	H.CONFIGURATION_NAME = 'name of model'; I.INPUT_NAME = 'channel name';

Replace the filters in quotation marks with the names for the database attributes.

SPSS view:

```
SPSSSCORE_V_LOG_HEADER AS H
Join
SPSSSCORE_V_LOG_INPUT AS I on H.SERIAL = I.SERIAL
```

Table 25. OFFER_MADE table for Predictive Customer Intelligence mapped to the custom database tables from IBM SPSS

Predictive Customer Intelligence Column	SPSS Column	Filters	Transformations
CHANNEL_ID	I.INPUT_VALUE	H.CONFIGURATION_NAME = 'name of model'; I.INPUT_NAME = 'channel cd';	NAME Lookup ID from CD
CAMPAIGN_ID	I.INPUT_VALUE	H.CONFIGURATION_NAME = 'name of model'; I.INPUT_NAME = 'campaign cd';	NAME Lookup ID from CD
OFFER_COUNT	Count (distinct H.STAMP)	H.CONFIGURATION_NAME = 'name of model'; I.INPUT_NAME = 'channel cd';	NAME
LOG_DATETIME	H.STAMP	H.CONFIGURATION_NAME = 'name of model'; I.INPUT_NAME = 'channel cd';	NAME

Replace the filters in quotation marks with the names for the database attributes.

SPSS View:

```
SPSSSCORE_V_LOG_HEADER AS h
join SPSSSCORE_V_LOG_OUTPUT on h.SERIAL = o.SERIAL
left outer join dbo.SPSSSCORE_V_LOG_INPUT li
on h.SERIAL = li.serial
```

Table 26. OFFER_RESPONSE table for Predictive Customer Intelligence

Predictive Customer Intelligence Column
CHANNEL_ID
CAMPAIGN_ID
RESPONSE_ID
OFFER_COUNT
LOG_DATETIME

You cannot get the customer response from IBM Analytical Decision Management. The customer response must be loaded from the channel application by using custom code.

Table 27. RESPONSE_TYPE table for Predictive Customer Intelligence mapped to the custom database tables from IBM SPSS

Predictive Customer Intelligence Column	SPSS Column	Filters
RESPONSE_TYPE_ID	Sequentially generated number	Distinct rows so Response Types do not repeat
RESPONSE_TYPE_CD	I.INPUT_VALUE	H.CONFIGURATION_NAME = 'name of model'; I.INPUT_NAME = 'response type cd';
RESPONSE_TYPE_NAME	I.INPUT_VALUE	H.CONFIGURATION_NAME = 'name of model'; I.INPUT_NAME = 'response type name';

IBM SPSS view:

```
SPSSSCORE_V_LOG_HEADER AS H
join SPSSSCORE_V_LOG_INPUT AS I on H.SERIAL = I.SERIAL
```

Appendix C. Troubleshooting a problem

Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and how to resolve the problem.

Review the following table to help you or customer support resolve a problem.

Table 28. Troubleshooting actions and descriptions

Actions	Description
A product fix might be available to resolve your problem.	Apply all known fix packs, or service levels, or program temporary fixes (PTF).
Look up error messages by selecting the product from the IBM Support Portal, and then typing the error message code into the Search support box (http://www.ibm.com/support/entry/portal/).	Error messages give important information to help you identify the component that is causing the problem.
Reproduce the problem to ensure that it is not just a simple error.	If samples are available with the product, you might try to reproduce the problem by using the sample data.
Ensure that the installation successfully finished.	The installation location must contain the appropriate file structure and the file permissions. For example, if the product requires write access to log files, ensure that the directory has the correct permission.
Review all relevant documentation, including release notes, technotes, and proven practices documentation.	Search the IBM Knowledge Center to determine whether your problem is known, has a workaround, or if it is already resolved and documented.
Review recent changes in your computing environment.	Sometimes installing new software might cause compatibility issues.

If the items in the table did not guide you to a resolution, you might need to collect diagnostic data. This data is necessary for an IBM technical-support representative to effectively troubleshoot and assist you in resolving the problem. You can also collect diagnostic data and analyze it yourself.

Troubleshooting resources

Troubleshooting resources are sources of information that can help you resolve a problem that you are having with an IBM product.

Support Portal

The IBM Support Portal is a unified, centralized view of all technical support tools and information for all IBM systems, software, and services.

The IBM Support Portal lets you access all the IBM support resources from one place. You can tailor the pages to focus on the information and resources that you need for problem prevention and faster problem resolution. Familiarize yourself

with the IBM Support Portal by viewing the demo videos (https://www.ibm.com/blogs/SPNA/entry/the_ibm_support_portal_videos).

Find the content that you need by selecting your products from the IBM Support Portal (<http://www.ibm.com/support/entry/portal>).

Before contacting IBM Support, you will need to collect diagnostic data (system information, symptoms, log files, traces, and so on) that is required to resolve a problem. Gathering this information will help to familiarize you with the troubleshooting process and save you time.

Service request

Service requests are also known as Problem Management Reports (PMRs). Several methods exist to submit diagnostic information to IBM Software Technical Support.

To open a PMR or to exchange information with technical support, view the IBM Software Support Exchanging information with Technical Support page (<http://www.ibm.com/software/support/exchangeinfo.html>).

Fix Central

Fix Central provides fixes and updates for your system's software, hardware, and operating system.

Use the pull-down menu to navigate to your product fixes on Fix Central (<http://www.ibm.com/systems/support/fixes/en/fixcentral/help/getstarted.html>). You may also want to view Fix Central help.

IBM developerWorks

IBM developerWorks[®] provides verified technical information in specific technology environments.

As a troubleshooting resource, developerWorks provides easy access to the most popular practices, in addition to videos and other information: developerWorks (<http://www.ibm.com/developerworks>).

IBM Redbooks

IBM Redbooks[®] are developed and published by the IBM International Technical Support Organization, the ITSO.

IBM Redbooks (<http://www.redbooks.ibm.com>) provide in-depth guidance about such topics as installation and configuration and solution implementation.

Software support and RSS feeds

IBM Software Support RSS feeds are a quick, easy, and lightweight format for monitoring new content added to websites.

After you download an RSS reader or browser plug-in, you can subscribe to IBM product feeds at IBM Software Support RSS feeds (<https://www.ibm.com/software/support/rss>).

Log files

Log files can help you troubleshoot problems by recording the activities that take place when you work with a product.

Error messages

The first indication of a problem is often an error message. Error messages contain information that can be helpful in determining the cause of a problem.

Notices

This information was developed for products and services offered worldwide.

This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service. This document may describe products, services, or features that are not included in the Program or license entitlement that you have purchased.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Software Group
Attention: Licensing
3755 Riverside Dr.
Ottawa, ON
K1V 1B7
Canada

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

Trademarks

IBM, the IBM logo and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “ Copyright and trademark information ” at www.ibm.com/legal/copytrade.shtml.

The following terms are trademarks or registered trademarks of other companies:

- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Index

A

accessibility for the Usage Report 63
Analytical Decision Management 8, 31
Apriori algorithm 34
Apriori model 45
Association model 10, 29
auto churn model 40
Auto Churn model 44
auto policy 44

B

banking 22
banking artifacts 70
banking offers 23
banking workflow 22
Bayesian Network algorithm 36, 50
Buying Behavior model 16
buying behavior segmentation 33
buying propensity model 40
Buying Propensity model 45

C

Campaign 7
campaign response model 40
Campaign Response model 44
case studies 3
Category Affinity model 16, 31, 39, 50
CHAID 44
CHAID algorithm 27, 42, 50
churn 27, 44, 50
Churn 10, 42
churn prediction model 40
Churn Propensity score 18
CLTV 20, 43
CLTV model 40
configure SPSS stream 58
Cox model 43
credit card debt 50
Credit Rating Model 49
Customer Affinity model 35
customer lifetime value 20
Customer Lifetime Value model 16
Customer Satisfaction model 28
customer segmentation 42
Customer Segmentation model 51

D

data 41
data model for the usage report 62
Data Processing Stream 45
data sources 25
Decision List algorithm 44
demand response program 21
demand response program assistance 49
demographic segmentation 33
deploy retail predictive models 37

deployment 52
Derive node 25

E

Energy and Utilities 21, 49
Energy and Utilities predictive models 48
Energy and Utilities sample 71
External Callout connector 18

F

Framework Manager model 63

G

generate streams 56

I

IBM Campaign 7
insurance 17
Insurance data model 46
Insurance Policy Recommendation model 45
inventory cost analysis 37
Inventory-based Suggestion model 13
Inventory-Based Suggestion model 37

K

K-Means clustering model 33

L

life policies 45
life stage segment model 40
Lifestage Segment model 45

M

Market Basket Analysis 39
Market Basket Analysis model 13, 34
master data 41
mobile banking 22
Mobile Banking Campaign 23

N

NBAOPT 55, 56, 58
Net Promoter Score model 28
Next Best Action Optimizer 55
Next Best Action Optimizer Studio 56

O

ODBC 25
offers 8

online behavior segmentation 33
online purchase funnel 13
optimization equation 14

P

PEV 46
PMML file 58
PMML files 56, 57
policy data 41
predict credit rating 49
Predictive Enterprise View 46
predictive models 10, 25
profile data 23

R

Recommended Rate Plan model 49
reports 61
Response Log Analysis model 36
Response Propensity model 30
retail case study 11
rules 8

S

samples 67
scoring a mode 51
segmentation model 40
Segmentation model 39
Select node 25
Self-Learning Response Model (SLRM) algorithm 36

sentiment 49
sentiment scoring model 40
Sentiment Scoring model 46
sequence analysis 51
SLRM 36
SMA Text Analytics model 46
social media 41
Social Media Analytics (SMA) model 40
SPSS Modeler 10
SPSS Modeler Text Analytics 20
SPSS stream 57
streams 56

T

Table node 25
Telecommunication sample 67
telecommunications 3, 4, 10
Text Analytics 46
Time Series model 37
transaction data 41
TwoStep cluster 42
TwoStep model 49

U

Usage Report 62, 63, 73

V

Variable File node 25