

Enabling End-User Datawarehouse Mining  
Contract No.: IST-1999-11993  
Deliverable No. D17.2

# Churn Analysis Case Study

Marco Richeldi and Alessandro Perrucci

Telecom Italia Lab  
Via G. Reiss Romoli 274 – 10148 Torino, Italy  
{marco.richeldi, alessandro.perrucci}@tilab.com

December 17, 2002

## **Abstract**

The case study that was developed to evaluate the MiningMart system is introduced.

A business model is provided and the pre-processing chains created by MiningMart are described.

# 1 Introduction

Telecom Italia Lab (TILAB) carried out extensive experiments to verify whether MiningMart project achieved its objectives. Project results were applied to perform a real-world discovery task that was accomplished by TILAB in the past by using a number of pre-processing and predictive modeling technologies.

The case study concerns developing a Churn Analysis system based upon data mining technology to analyze the customer database of a telecommunication company and predict customer turnaround.

Case study business model is introduced in Chapter 2. Pre-processing chains are described in detail in Chapter 3.

## 2 The Case Study – Business model

Evaluation of MiningMart was performed by carrying out a difficult knowledge and data discovery task concerning understanding the behavior of customers of a telecommunication company that unsubscribe and move their business to a competitor and predicting the likelihood of this event to occur. Customers who turnaround are called churners and the process of exploring and modeling their behavior is named Churn Analysis.

The Churn Analysis case study was selected as the massive volume and complexity of data collected by telcos about their customers and the difficulty of modeling this data make churn analysis a very good test bed for evaluating MiningMart features. In addition, TILAB developed several Churn Analysis systems for different companies of the Telecom Italia Group. The knowledge and data discovery process required to build effective churn prediction models has been widely explored along the years, and it reveals quite easy to assess how MiningMart supports it.

### 2.1 Churn Analysis

Customers become “churners” when they discontinue their subscription and move their business to a competitor. That is, churning is the process of customer turnover. This is a major concern for companies with many customers who can easily switch to other competitors. Examples include credit card issuers, insurance companies and telecommunication companies.

The telecommunication industry has seen many changes in the few last years. Examples are liberalization of markets – almost all European countries are now allowing foreign carriers to compete on local and long distance services –, de-regulation and a fierce competition from other industry sectors. An immediate manifestation of the new competitive environment is the escalating problem of churn. Long-distance carriers and cellular companies are well acquainted with the issue of valued customers unsubscribing to purchase product and services offered by a competitor. In the United States, an estimated two millions long-distance customers churn each month. The monthly churn rate for European cellular carriers averages between 8 and 12 percent.

The cost of churn in the telecommunication industry is large. The cost of churn in wireless communication may be around 500 euros. In addition, it costs a great deal more to win new customers than it does to retain current ones. And, frequently, a new

customer will churn away before the company can fully recoup its acquisition costs. It is clear that spending money holding on to existing customers is more efficient than acquiring new customers.

Consequently, churn management has emerged as a crucial competitive weapon, and a foundation for an entire range of customer-focused marketing efforts. With effective churn management, a company is able to determine what kind of customers are most likely to churn, and which ones are most likely to remain loyal. Part of the process is determining customer value, as sometimes you may want to let go customers that are not or very little profitable. When this kind of knowledge is available to a company, marketing managers may take informed and strategic action to minimize defections, win back valued defectors, and attract more cost-effectively the right kind of customers in the future – including those that are least likely to churn.

However, given that most customers will only signal their intention to churn when they call to cancel their account, it is difficult using standard techniques to target for anti-churn marketing. One solution to combating churn in telecommunications industries is to use data mining techniques. Data mining may be used in churn analysis to perform two key tasks:

- Predict whether a particular customer will churn and when it will happen;
- Understand why particular customers churn.

These two tasks are referred to as “prediction” and “understanding”. They represent the two most important aspects of data mining in use today. By predicting which customers are likely to churn, the telecommunication company can reduce the rate of churn by offering the customer new incentives to stay. By understanding why customers churn the telco can also work on changing their service so as to satisfy these customers ahead of time. In addition, the chance of the customer churning after action is taken can be assessed by the data mining tool so as to choose the best strategy in terms of cost and effort.

Before either prediction or understanding are performed, however, data must be gathered. This is typically done in a data warehouse, which is a large repository of clean, non-volatile and historical data. The types of information that are stored in a data warehouse for a churn analysis include:

- Customer demographics, i.e., age, gender, marital status, location, etc.
- Call statistics: length of calls at different times of the day, number of long distance and local calls.
- Billing information for each customer – what the customer is paying for local and long distance.
- Extra service information, that is, what extra plan the customer is registered on, e.g. special long distance rates.
- Voice and data product and services purchased by the customer, e.g., broadband services, private virtual networks, dedicated data transport links, etc.
- Complaint information: how many customer service calls are made for disputed billing, dropped calls, slow service provisioning, unworking special services, and so on.
- Credit history.

## 3 Case study description

### 3.1 Input data

In this case study we describe a data mining solution to predict and understand churn behavior of customers of a mobile operator. The raw data files were generated by different Operational and Business Support Systems and are structured into the following large tables, namely:

- Customers: basic customer information;
- Cdr: call data records aggregated by month;
- Services: basic information concerning service subscribed;
- Revenues: basic billing data aggregated by month.

The Figure 1 depicts source tables and their attributes and relationships.

**CUSTOMERS** table: contains customer information: a key that identifies customers in the source systems and other basic information, e.g., age and gender.

**CDR\_T** table: contains call data records (CDRs) in transactional form. Each record provides the duration of a specific type of calls (peak, off peak, weekend, international calls) aggregated on a monthly basis (5 months are available: month labels are from 1 to 5).

**SERVICES** table: contains basic information concerning service subscribed by customers, such as: handset (cell-phone model/class), length of service (contract duration in months), number of dropped calls, tariff plan and tariff type used for billing, and target attribute, that is, the churn state in month 6.

**REVENUES** table: contains basic billing data in transactional form. Each record provides the revenue (difference between income and costs) related to each user aggregated on a monthly basis (from month 1 to month 5).

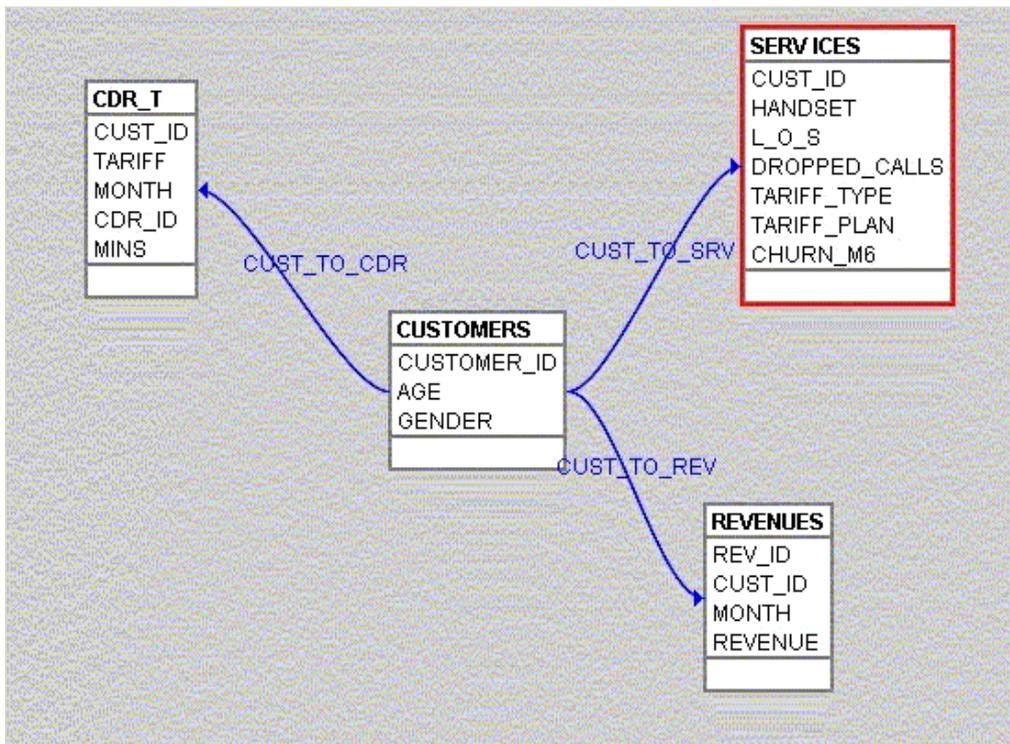


Figure 1: Conceptual data model - source tables and their attributes and relationships

## 3.2 Discovering tasks

### 3.2.1 Pre-processing chains

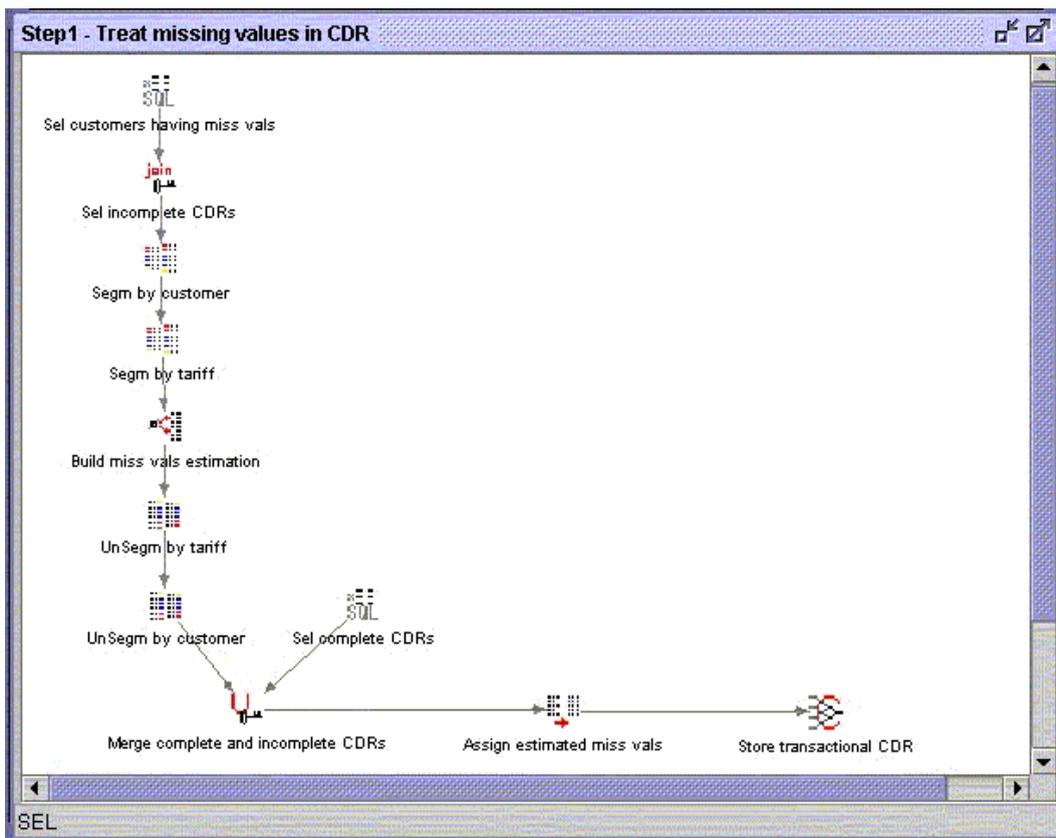
The data mining process has been divided into five tasks as follows:

- Handling missing values in CDR;
- Transpose CDR from transactional form to relational form;
- Transpose REVENUES from transactional to relational form;
- Create derived attributes and customer profile;
- Churn prediction modeling;

For each task a folder chain containing one or more operator chains has been developed.

### 3.2.2 Handle missing values in CDRs

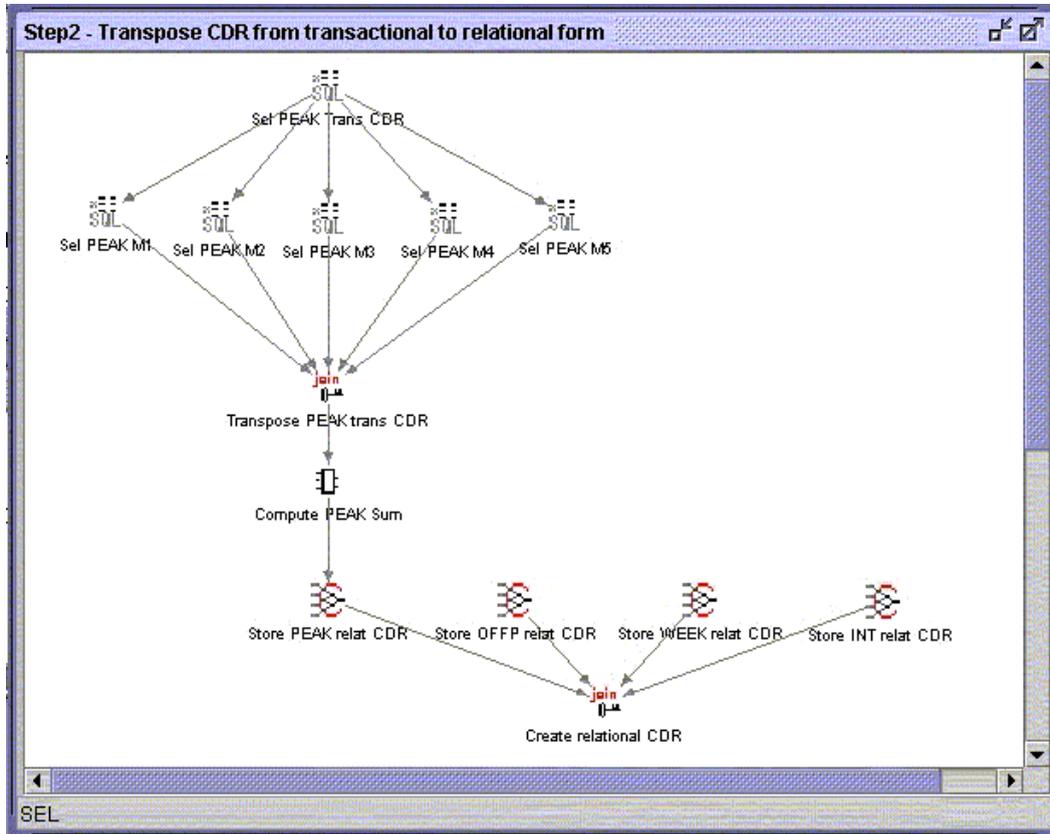
This task performs missing values estimation in call data records (CDRs) represented in transactional form.



- **Sel customers having miss vals:** Filter out all those customer for whom at least one CDR in the set of CDRs generated by the customer features a missing value. A list of keys of customers with missing values in associated CDRs is generated.
- **Sel incomplete CDR:** Select incomplete CDRs (CDRs with missing values) by joining the list of customers keys generated by the “Sel customers having miss vals node” and the CDR table.
- **Segm by customer:** Create a view containing incomplete CDRs for each customer.
- **Segm by tariff:** Create a view containing incomplete CDRs for each tariff and customer.
- **Build miss vals estimation:** Substitute any missing value in a CDR with the average of the attribute values available in the CDR. Adds the result to the CDR view.
- **UnSegm by tariff:** Rebuild incomplete CDR views for each tariff and customer.
- **UnSegm by customer:** Rebuild incomplete CDR views for each customer. The result consists of incomplete CDRs where missing values have been filled in using the estimate provided by “Build miss vals estimation” node.
- **Sel complete CDR:** Select CDRs with no missing values.
- **Merge complete and incomplete CDRs:** Merge complete and incomplete CDRs (by substituting missing values with their estimates).
- **Assign estimated miss vals:** Assign missing values their estimation.

- **Store transactional CDR:** Store transactional CDRs by joining CDRs and customer records by customer key. Discard attributes used to perform the missing values estimation process.

### 3.2.3 Transpose CDR from transactional to relational form



This task applies a pivoting operator that performs a transposal operation to convert call data records (CDRs) stored in transactional form into CDRs stored in relational one.

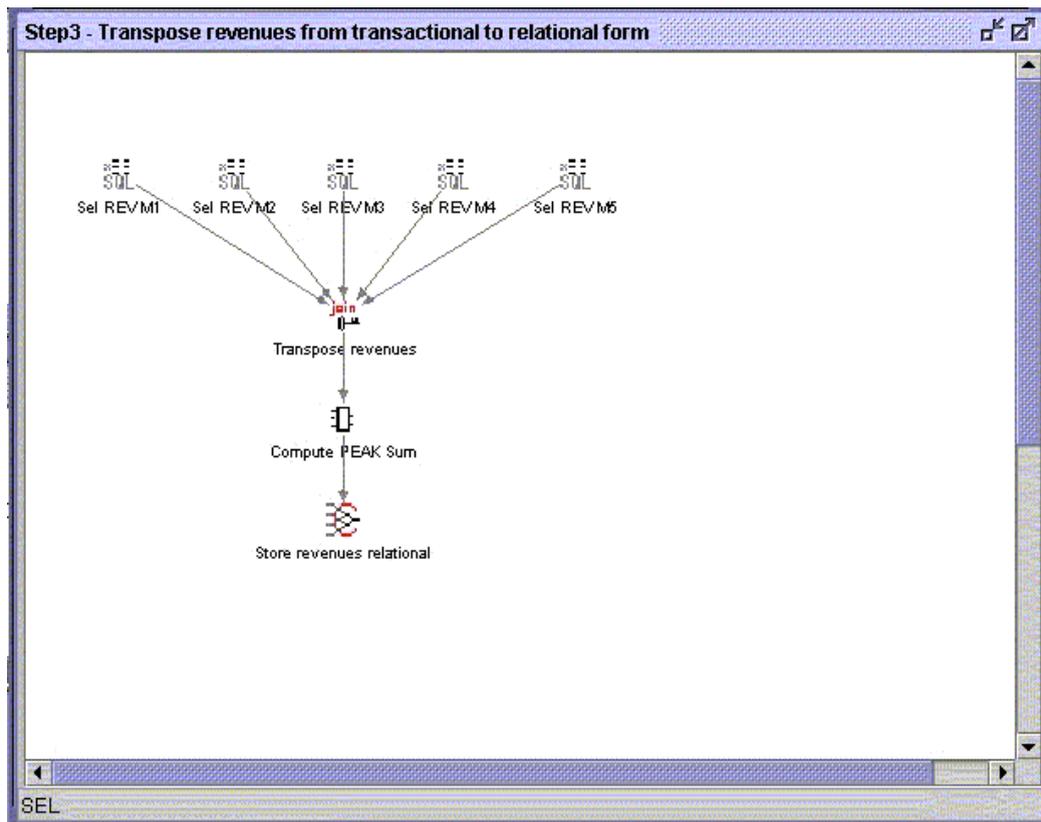
In transactional form a CDR includes the month when it was generated (CDR have been previously aggregated on a monthly basis) and the length for a call of given type. Consequently, many CDR may be available that describe the call behavior of a customer within the observed time window of 5 months (one CDR for every month for every call type). In relational form, just one CDR is required to describe the call behavior of the customer, since call lengths made by the customer in any month for any call type are included in a single data record.

Different type of calls (i.e., PEAK, OFFP, WEEK and INT) are processed by different chains. The final step of any of these chains joins together CDRs of different type and generate a data set in relational form.

Each chain includes the following five steps: selecting CDRs by tariff, selecting CDRs by month, transposing CDRs into relational form, calculating derived attributes and eventually saving the new enriched CDRs.

- **Sel PEAK trans CDR:** Select transactional CDRs associated with calls of PEAK type.
- **Sel PEAK M1 ... Sel PEAK M5:** Select CDRs associated with calls of PEAK type performed in a specific month (from M1 to M5).
- **Transpose PEAK trans CDR:** Convert CDRs associated with calls of PEAK type from the transactional form to the relational one. Indeed, it first joins together all selections generated by "Sel PEAK M1 ... Sel PEAK M5" node and then renames CDRs so that the new name recalls the call type (e.g., PEAK) and the month when the call was generated (labels from M1 to M5).
- **Compute PEAK Sum:** Add to CDRs associated with calls of PEAK type a new attribute that sums up the duration of all calls performed from month M1 to month M5.
- **Store PEAK relat CDR:** Save CDRs associated with calls of PEAK type by joining CDRs and customer records by customer key. It discards attributes that were used to support the pivoting operation.
- **Create relational CDR:** Join together all CDRs associated with different call types, i.e., PEAK, OFFP, WEEK call types.

### 3.2.4 Transpose REVENUES from transactional to relational form

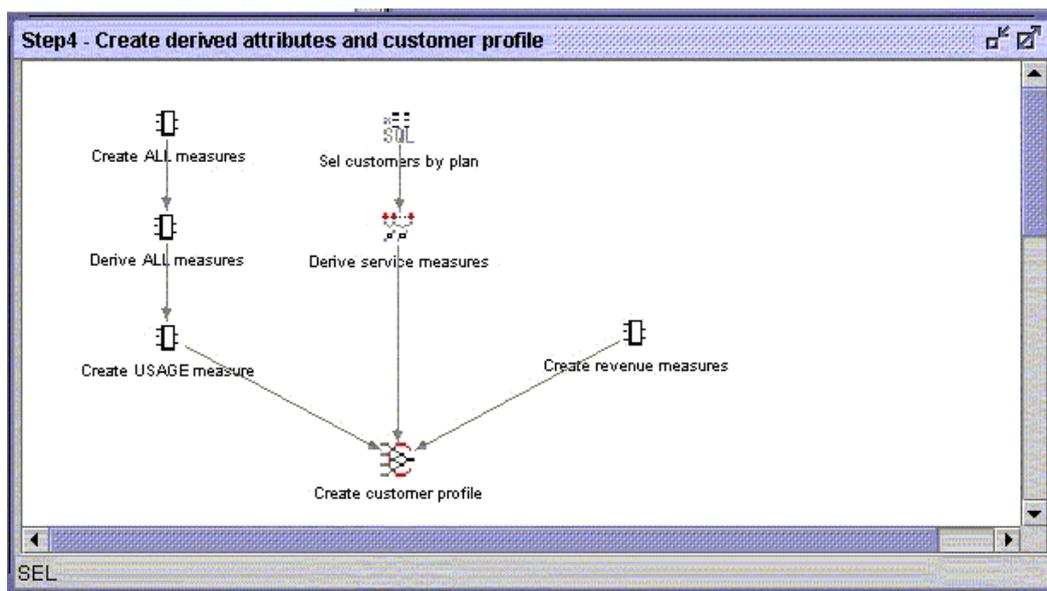


This task applies a pivoting operator that performs a transposal operation to convert REVENUES records stored in transactional form into REVENUES records stored in relational one.

The pre-processing chain transposes revenue records from the transactional form into the relational one. Four processing steps are required: selecting revenue records by month, converting revenue records into relational form, calculating derived attribute and eventually saving the new revenue records.

- **Sel REV M1 ... Sel REV M5:** Select revenue records associated with calls originated in a given month (from M1 to M5).
- **Transpose REV trans CDR:** Convert revenue records from a transactional form into a relational one. This node joins together revenue records selected by the “Sel REV M1 ... Sel REV M5” node. Every selected revenue attribute is renamed by appending a suffix that provides the month when the call was originated to a prefix (namely, REV) that reminds that it is a revenue record.
- **Compute REV Sum:** Add a new attribute that sums up the revenue of calls originated from month M1 to month M5.
- **Store REV relat CDR:** Save revenue records by joining revenue records in relational form and customer records by customer key. Attributes used to perform the pivoting operation are discarded.

### 3.2.5 Create derived attributes and customer profile

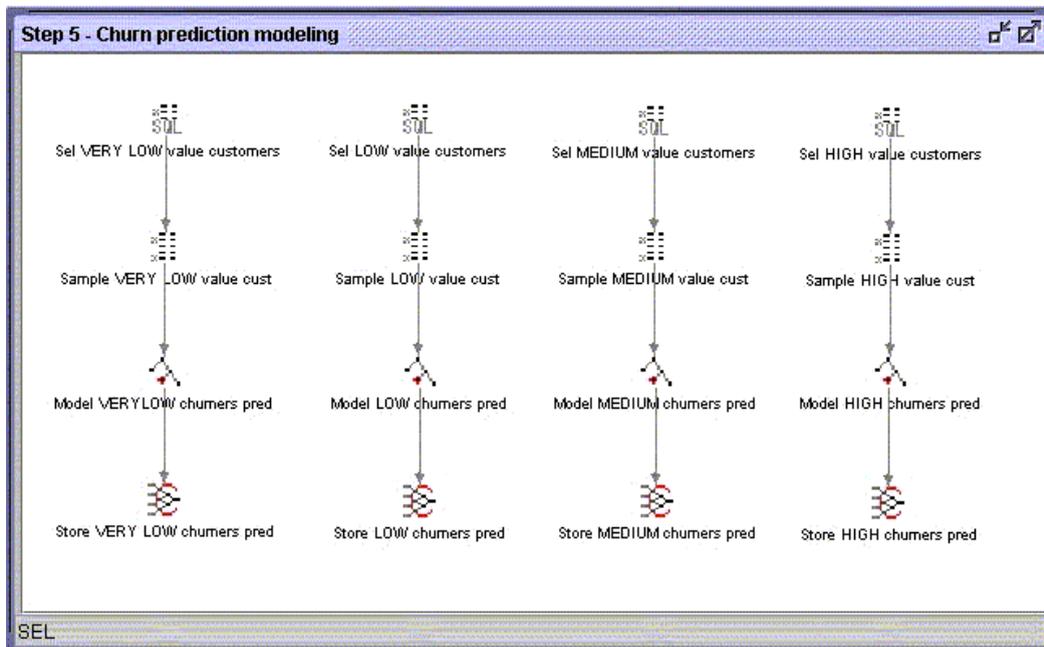


This task extends information power of source data. Construction operators are applied to create new attributes that provide interesting knowledge about the customer behavior at a higher level of detail with respect to CDRs. Derived attributes are joined together in a new customer profile table. The pre-processing chain includes the following nodes:

- **Create ALL measures** Calculate call duration by aggregating CDRs on a monthly basis. Adds duration of peak, off peak, week and international calls for every single month – from month M1 to M5.
- **Derive ALL measures** Calculate call duration at the month level of aggregation.

- **Create USAGE measure:** Calculate difference between call durations for different time lags. That is, four differences (namely ALL\_DV51, ALL\_DV52, ALL\_DV53, ALL\_DV54) are computed by subtracting call duration at months M1, M2, M3 and M4, respectively, from call duration at month M5. Further, slope of the straight line (ALL\_ANG51) that passes through the points whose Y coordinates are the call durations at month M1 and M5 is computed.
- **Sel Customers by plan:** Selects customers by tariff plan. There exists 3 tariff types (CAT, PLUS and PLAY). Investigation was focussed on CAT and PLAY plans as the most interesting ones.
- **Derive service measures** Apply a discretization operator to dropped call attribute DROPPED\_CALLS and length of service attribute L\_O\_S. Q\_O\_S BAND and L\_O\_S band attributes are derived.
- **Create revenue:** Apply a discretization operator to the attribute providing overall revenue by customer (REV\_TOT\_SUM). VALUE\_BAND attribute is derived.
- **Create customer profile:** Join new attributes created in the previous steps.

### 3.2.6 Churn modeling chain



This task is aimed at building a predictive model for each of the four groups of customers originated by segmenting customers according to their value (basically, customers are grouped according to the value of attribute VALUE\_BAND).

Four similar processing chains have been defined. Every chain includes the following steps:

- **Sel HIGH value customers:** High value customers are selected.
- **Create a random sample:** A random sample is drawn to create a training set.
- **model HIGH churners pred:** The decision tree operator is applied to fit the data in the training set and create a model that predicts the likelihood of a customer to become a churner in the month M6.

- **Store HIGH churners pred:** Save output of the modeling stage in a view table, so that performance evaluation of the final model can be carried out.