

Championing of an LTV Model at LTC

Edmund Freeman
Washington Mutual Bank
1201 3rd Avenue
Seattle, WA 98101
(206) 461-7883

edmund.freeman@wamu.net

Gabor Melli
PredictionWorks Inc.
Seattle, WA
(206) 369-3582

gmelli@predictionworks.com

ABSTRACT

In this paper we report on the successful implementation of a life-time value (LTV) forecasting system at a large telecommunications company (LTC). While some research results have been reported elsewhere on the technical challenges of modeling customer value, our experience suggests that a data mining system implementation can expect to encounter several organizational challenges that can impinge on its success. We provide a background on the application, and then analyze several success factors.

Categories and Subject Descriptors

H.2.8 [Database Applications]: Data Mining

General Terms

Management, Economics, Human Factors.

Keywords

Organizational Change, Data Mining, Customer Lifetime Value, Successful Data Mining.

1. INTRODUCTION

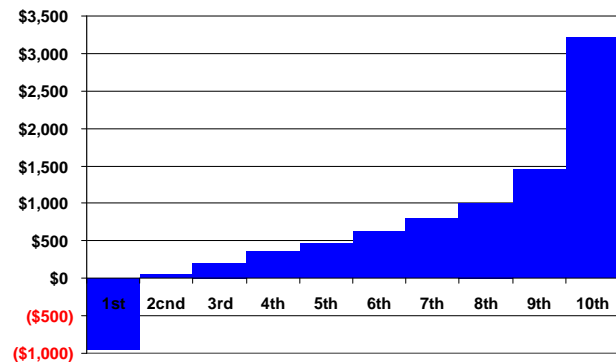
One of the application areas that data mining has begun to assist in industry is in the forecasting of customer value [3,6]. The application is particularly useful in the service sector where a large variance exists between the revenue generated by a customer and the cost associated with delivery of the service to that customer [4]. The differential between these two monetary figures, over a customer's entire relationship with the organization, is often referred to as the customer's life-time value (LTV, or sometimes as CLV). Figure 1 illustrates the breadth of the range in customer value at a telecommunications company. Noticed that a small but substantial portion of the customers have negative value. This situation can occur for example when a customer makes extensive use of an unprofitable feature that exists in order to match a competitor's offering. If an organization could predict what customers will continually generate negative value then it could respond by avoiding any optional costs on these customer. The organization could for example, avoid the cost of mailing promotions to these customers.

In 2002, we participated in just such a project at a large telecommunications company (LTC). LTC was a telecommunications company with approximately 18M customers and \$12B a year in revenue. Because much of their revenue came from a subscription-based service, one of LTC's larger internal groups was its customer relationship management (CRM) division. The goal of CRM is to increase revenue and customer

satisfaction by keeping customers informed about services and offers that will appeal to them. Effective communication however has an incremental cost. It would be helpful if the company could predict which to which customer it would make economic sense to contact.

Figure 1 –Distribution of customers by their future value, divided into ten equal sized groups (deciles). The negative-value customers account for approximately 8% of the customer base.

Forecast LTV Scores by Decile



Based on a preliminary analysis we estimated that implementing the tactic of not promoting negative-value customers would, for a given marketing campaign, decrease campaign-related expenses by 5% while increasing the campaign returns by 10%. Figure 1 illustrates this segment of the customer base. The analysis also demonstrated that it was feasible to predict life-time value with sufficient accuracy. The remaining question was the cost of implementation. The estimation of this cost was for approximately 0.1% of CRM's annual budget. The system would pay for itself within two months, and generate a 500% ROI.

Easier said than done. Despite the significant opportunity several critical organizational challenges were encountered both during the implementation and soon after its launch. The project required many departments to re-think LTC's relationship with their customers and the departmental relationships within LTC. Based on past experience, we suspect that any successful LTV project will impel such an organizational reevaluation.

The first part of this paper presents an overview of the technical challenge that the implementation team faced. The remainder of the paper presents on the less well reported aspect of maneuvering the successful implementation of such a project through the organization.

2. LTV SYSTEM DESIGN

The requirements of the LTV system were to produce a set of values and scores for each customer on a monthly basis. The main outputs were the following:

Forecasted Value: The remaining monetary value predicted for a customer. For example, a forecast value \$3,721 for customer x would be an accurate forecast if when this customer terminated their account the account would have resulted in an additional \$3,721 in value.

$$Forecast = \sum_{t=1}^{t=60} M(1-v)^t(1-i)^t / (1+r)^t$$

Where M is the average profitability over the past 12 months, v is the probability of voluntary attrition per month, i is the probability of involuntary attrition per month, and r is the Net Present Value rate.

Past Value: The profitability to-date, including acquisition cost, for a customer. For example, a past value of (minus) -\$73 for customer y signals has not yet resulted in a profit. LTC had substantial acquisition costs, in the \$350 - \$550 range per customer; it took on average two years for a customer to become profitable. Past Value was measured by keeping a running sum of monthly profitability minus the acquisition costs, and acquisition costs were measured by acquisition channel and quarter.

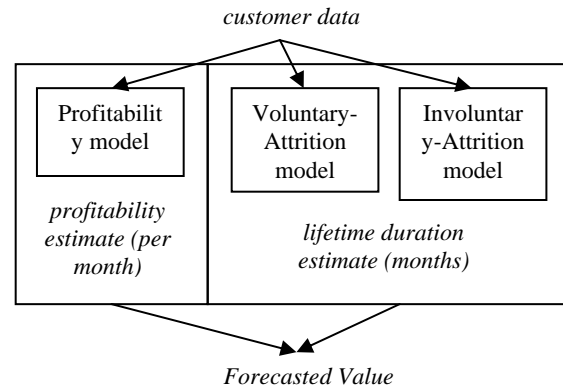
Expected Lifetime Value: The total value expected for this customer. This value is simply the summation of forecasted value and past value.

Potential Value: Another derived measure that proved to be useful was the Forecasted value with no attrition. Because Potential Value does not involve attrition, it was easy to overlay attrition data onto the Potential Value and create a cluster of high-potential, high-attrition customers to focus retention efforts on.

2.1 LTV Model

The approach taken to modeling LTV was to divide the task into three separate optimization problems. The first task was to model the customer's average expected monthly profitability. The second and third optimization tasks modeled different scenarios of when the customer would cease to be a customer. The first scenario modeled was when the customer makes the request that their service be terminated. In the industry this event is referred to as voluntary churn or voluntary attrition. The second scenario modeled was when the company requires that a customer's service be terminated. This is referred to as involuntary churn (or involuntary attrition), and is typically the result of non-payment. Together, along with a Net Present Value factor, the three models were combined to calculate life-time value [Figure2].

Figure 2 – Data flow for the calculation of the LTV Forecasted Value.



Other more sophisticated approaches were considered. For the attrition models a survival model such as the Cox proportional hazards model [1] could have been a relevant methodology to attempt.

The simplicity of the three sub-model approach did have a substantial benefit after the project was completed in that it was simple to modify the scores to meet special needs. For instance, 'what-if' scores were quickly introduced simply by removing off-network expenses and bad debt expenses so that CRM analysts could see the value of addressing those issues.

2.2 Definition of Profitability

The defining of profitability was an important preliminary task. Mixed advice was encountered about whether to use either the expenses in the financial statements or of releasing a general survey on expenses. Using hard financial statements made the LTV formulas solid and allowed freedom from organizational misperceptions about customer value. For example, off-network and bad debt were two of the more critical expenses to LTV calculation, but few departments were interested or understood these expenses. Instead the expenses of customer care and promotions, that had relatively minor effect on LTV, would have been overrepresented because of their organizational wide visibility.

After the expense categories were established a decision was made on (1) how to best allocate that category to customer activity and (2) whether the category should be included at all.

Expense	Allocation	Inclusion
Off-Network Time	Minutes off network multiplied by the fee charged for the individual network used.	Yes
On-Network Time	Minutes on network multiplied by network maintenance charges.	Yes
Long-Distance	Minutes of long distance usage times a per-minute charge.	Yes
Customer Care Calls	Number of calls to care multiplied by a cost per call.	Yes
Bad Debt Expense	Revenue weighted by probability of default.	Yes
Misc. System	Percent of gross revenue.	Yes.

Expense	Allocation	Inclusion
Expense		
Misc. Expenses	Flat amount per customer.	Yes
Network Capital Expense	Minutes on network multiplied by network capital expenses.	No!

The handling the capital expense was one of the more critical challenges encountered. The Finance division wanted the capital expense included because it is part of their LTV calculation. The addition of this term however would be problematic because, as we illustrate below, it would result in 25% of the customer base as having a negative value, instead of the more reasonable 8%. This change would have resulted in substantial operational difficulties for us.

In one sense either of these two versions of zero (0) LTV would be acceptable. However, when presented with a profitability analysis the natural inclination for a business manager is to do the things that are profitable and to avoid those that are unprofitable. A system would like be more successful if it worked with this human inclination and not against it. Before implementation we delved into a deeper understanding of the effects of identifying portions of our customer base as negative-valued.

2.3 Profitability Death Spiral

One of the lessons for future projects is that the intuitive application of data mining scores can lead to undesirable consequences. A nice example comes from a manufacturing setting [personal communication]. The company in question ran several plants at over-capacity and other plants at 60% capacity. It also possessed a relatively accurate profitability model. In the costing model however, capital expenses were allocated by unit produced. The over-capacity plants had a much lower unit cost than the plants running at 60%. The reason for this outcome was that when the profitability data was first published there were minor variations in production, which in turn resulted in minor variations in per-unit profitability. Naturally the company increased production in the more-profitable plants and decreased production in the less-profitable plants. Over time the random variations were amplified and resulted in the inefficient situation.

Because LTC decisions were typically measured by their ability to reduce attrition, we estimated the likely effect on attrition of removing 8% and 25% of the population from the two major campaigns, contract renewal and equipment upgrade. Approximately 60,000 customers responded to the contract renewal program each month, and 42,000 customers a month received an equipment upgrade. Then each month we would have

	8% Negative	25% Negative
Contract Renewal	60,000	60,000
Effectuated	4,800	15,000
Attrition Rate	16%	16%
<i>Extra Attrition</i>	768	2,400
Equipment Upgrade	42,000	42,000
Effectuated	3,360	10,500
Attrition Rate	35%	35%
<i>Extra Attrition</i>	1,176	3,675

<i>Total Extra Attrition</i>	1,944	6,075
<i>Attrition Increase</i>	1 basis points	4 basis points

Neither formula would cause unmanageable problems with customer attrition. However, there were two powerful operational reasons to chose the 8% negative solution and not include capital expense in the formulas. First, LTC was constantly managing to attrition and having fairly reliable attrition crises. At 25% negative it becomes a persuasive argument that we should ‘turn off’ LTV and make save offers regardless of value. At 8% negative it becomes much more reasonable to craft attrition solutions within the LTV system. Secondly was the type of customers that were identified as negative. Without capital expenses each customer that was negative had a clear profitability-destroying behavior. With capital expenses a large class of negative-valued customers were those that were simply using most of their plan minutes, and our business partners were very reluctant to negatively impact those customers.

2.4 Features of the Profitability Calculation

Our profitability formula had a number of important features that were critical to the success of the project.

- 1) The negative-valued customers were justifiably negative-valued. For each such customer, we could identify concrete behaviors as to why they were negative-valued and that LTC would in fact be better off without that customer. As can be seen in Figure 1, the negative-valued decile is clearly negative-valued.
- 2) For each critical component of profitability we could give concrete advice on how to improve it. For instance, for Bad Debt expense we could give suggestions on how to acquire more credit-worthy customers.
- 3) The formula was based on actual financial data, so we could make meaningful comparisons between customer value and marketing offers.

3. PROJECT PRECURSORS

This was not the first attempt at LTC for an LTV system. Two earlier attempts did not achieve a return on investment. One was technically successful but achieved minimal impact on the business; the other did not get past the proposal stage.

3.1 Finance: Ignored Valuation

The LTC Finance department produced a customer-based valuation. They published this information by rolling all the information up to the rate plan level and then producing profitability numbers. This information was ignored outside of Finance. This was because

- 1) Profitability by rate plan was not the perspective taken by others in the organizations.
- 2) The report was painfully dense; LTC had over 1,000 active rate plans.
- 3) There was no way to drill down into the data and identify causes of profitability and unprofitability.
- 4) The calculations used base averages and were not adjusted for customer behavior such as different attrition rates and non-payment rates. For instance, Bad Debt was allocated as a percent of gross revenue. Telling people to decrease bad debt by decreasing revenue is not very actionable advice.
- 5) Finance made no effort to get their results out into the company and have it be used.

This first attempt proved helpful later to better understand the financial dimension to calculating life-time value.

3.2 The Consultant-Lead Committee

Another attempt at an LTV system was a consultant-lead committee (CLC). Its approach to implementing an LTV system was to interview a large number of business-oriented subject-matter experts about all aspects of customer profitability. The end result was a long report that was soon shelved. This was likely because:

- 1) The proposal contained hundreds of recommendations that, while grounded on the experience of subject-matter experts, needed to be pared down to a more cost-effective subset of requirements.
- 2) The recommendations were not grounded on a theory of customer value. Instead the metrics were based on subjective opinion and as a result was unable to defend or explain its rationale.
- 3) Most critically, the CLC did not have the technical expertise to implement their LTV system. This is the real reason the project never got beyond presentationware.

One source of value from this attempt at LTV was ideas on how to present LTV results.

3.3 Precursor Lessons

Both previous projects ultimately failed because they did not result in data that was useable to the enterprise. If we wanted our project to be successful we needed to make the data available, which meant we needed to get the LTV scores into the data warehouse, which meant we needed IT funding, which meant we needed to go through LTC's new funding process.

4. IN ORDER TO GO LIVE

The technical aspects of the project took the team approximately one third of the year that it took to launch the LTV system. The rest of the time was meeting with partners, discovering who we did not need to meet with, getting support from the people we needed, going through IT, educating the company on what LTV was and how to use it, and building ancillary systems.

The LTV project brought the project team into conflict with

- New cost control methodologies in the company,
- Finance, in terms of how marketing campaigns were evaluated,
- More Finance, in terms of how corporate profitability was managed,
- IT, in terms of how data was managed and produced,
- A Small Influential Business Unit (SIBU), and
- Regional marketing units.

Why go through this risky effort when the LTV scores could have been generated on a high-end PC? Because the project would have been unsuccessful. A successful project required that the data be housed in the company data warehouse because this made the data universally accessible and also gave the LTV project a kind of official stamp of approval.

4.1 The Start: The New Cost-Control Process

The project's initial challenge was the process for controlling operational costs, specifically the information technology budget. The process for funding new projects called for strict return on investment (ROI) calculations, with a senior manager held

accountable for delivering the ROI, and the process validated by a committee selected from all units in LTC.

An ROI requirement is commonplace and reasonable but not necessarily a rational exercise. The committee members for example only had sufficient time for a cursory glance to the proposals. A consequence was that each member responded to the marching order of "get their department's projects funded", so project funding was less a matter of project merit than political connections and a friendly accountant who would give a favorable ROI evaluation.

Estimating project return presented us with a substantial difficulty. Marketing campaigns had been measured on either 1) attrition improvements or 2) new purchases. Gradually improving the behavior of the base was not in their formulas. The Finance customer valuation method allocated most costs on a per-customer basis, ignoring facts such as bad debt expense tends to be highly concentrated on customers who stop paying.

This was one of the points in the process when the use of a dedicated PC to produce LTV reporting was considered. However, the Marketing Vice-President (MVP) insisted that no such skunkworks project be done, and that we needed official funding.

In the end, the MVP figured how much ROI would be necessary to get the project through, what kind of attrition reduction would be necessary to get over that hurdle, and then promising to delivery that reduction. This was a complete prevarication; the LTV project was designed to vastly improve customer profitability at the cost of a slight increase in attrition. However, it was enough to get the project funded, at which point we ran into the Finance and IT department challenges.

4.2 Finance Department

Finance thought that an LTV project was a great idea. However, they were upset that Marketing was doing it and not them. They were already publishing a form of customer valuation, but because the information was not actionable it was not used. The situation between the two groups was understandably tense. The situation was made unavoidable by our executive sponsor's insistence that we have Finance's formal approval of our methods. Some unexpected challenges came from the day to day interactions with their team. Often these interactions involved phone conferences, but habitually meeting invitations would be unacknowledged. On the other hand, our Finance partners had a habit of showing up to meetings they were not invited to, so we had to be ready to discuss LTV at any time. Straight answers were also not always simple to come by. For example, once to the question "Is this how we should be handling this expense category?" the reply was "What would happen if we lost all our customers?"

In the end the process resembled a Poisson process with a low probability of success. It was simply a matter of trying again and again until they (somewhat accidentally) said yes. As we found out later, Finance finally approved our formula because they did not think the LTV project would actually get finished and that if the LTV project were to be finished it would not be taken seriously.

The long process of working with Finance did have beneficial results. We had to think very carefully about how we were evaluating customers, and we had a much more robust formula at the end. From our experience and personal communication, any

LTV project will require close and often contentious work with Finance.

4.3 Information Technology Department

A commonplace challenge to data mining projects is in the interaction with the information technology department. At LTC IT department's motto was "we will get you anything you want, just tell us what columns you want in your flat-file extract". Having another department producing production programming that would affect the data warehouse was a new idea to them.

We spent several months discussing the protocols of us working closely with IT programmers and establishing project resources (which included a very small disk space requirements on a server with spare capacity), only to have the whole plan shot down by IT management. The reason given was that the systems programming involved a model, and the IT department was not capable of handling models – only Marketing was. The result was that IT would drop off a data file and later pick up another file to load, but we would have to do all the programming in between.

At the end, this was a beneficial result, giving us necessary control over the output. However, the route did seem unnecessarily unpleasant. Managing IT's issues was primarily a matter of patience and flexibility.

4.4 "We Have to Stop This!": Small Influential Business Unit (SIBU)

SIBU was one of the groups we needed buy-in from. SIBU was responsible for a potentially highly profitable future line of business, and had identified a small group of current customers that they thought would be good targets for the new services. Because they were expected to be highly profitable in the future, SIBU had tremendous influence within LTC and their buy-off was needed for major projects.

SIBU's initial reaction to the project was absolute horror. Some of "their" customers might get poor scores! First, SIBU insisted that none of their potential customers get scored at all. They demurred when we pointed out this would mean essentially excluding them from all regular marketing efforts. SIBU's next idea was to stop the project completely.

We quickly realized that if we did not get the problem solved right then and there, the project would be dead. The solution to the challenge was the addition of a set of scores (A/B/C/D/E) that was based on a special usage formula that SIBU provided.

As it turned out, by the time LTV went into production business had changed and the SIBU scores were irrelevant; all customers received a 'C'. However, the scoring provided us the buy-in from SIBU necessary to complete the project without changing the core system.

4.5 Validating LTV

Ideally we would have validated the LTV scores by going back five years, scoring the customer base then, and then checking to compare our predictions versus reality. This was not possible. At LTC we only had 13 months of history available to us. In particular, we would not be able to test the critical assumption of the LTV system, that after an initial adjustment period a customer's current behavior would on average be the customer's future behavior.

What we did instead was to validate the decisions that the LTV system was indicating we should make. For example, the LTV

system was indicating that a substantial population of customers was unprofitable because of off-network charges. LTC had to pay competing telecommunications companies fees to support the usage of LTC customers. To validate the LTV system we got the individual customer's monthly usage, matched with our contract data to calculate how much each month we paid in fees to our competition, and could identify that we were losing substantial amounts of money keeping these customers as customers. This justified programs that gave substantial bonuses to customers to move to a plan where they paid for off-network usage.

There was yet another LTV project at LTC that failed this validation test. A European Consulting Company (ECC) had been brought in to do LTV in parallel with us. The ECC system relied heavily on network charges to calculate the monthly profit/loss on customers. Like our LTV system had identified off-network usage as an issue the ECC system identified large on-network usage as an issue, and ECC crafted a program to force migrate (i.e., 'fire') customers with high usage on family plans.

As the program managers drilled down into the data they became very dubious about the ECC program. They were targeting customers that were only using 50% of their allotted minutes, and this hardly seemed like excessive usage to the program managers. We found out about the ECC program when the program managers came to us for advice on the LTV scores of the targeted customers. The LTV scores showed no real difference between the targeted customers and other customers and this information was used to kill the program. This was the only time in my professional experience that I have seen program managers happy to have data kill a program.

5. AFTER PRODUCTION

Once the LTV system went into production, a new set of issues arose.

5.1 Customer Care

The first challenge after production was an unexpected demand to explain of individual scores. Fortunately the simple, modular nature of our LTV formula enabled quick and believable answers to all of these questions. For example the question "Why does Mr. Jones have high revenue and such a bad LTV score?" was commonly answered with "Mr. Jones had not paid us in X months". (Surprisingly, the customer care system did not take payment history into account when handing out equipment). The ability to quickly provide clear, convincing answers to valuation questions gave the project a tremendous amount of credibility in the enterprise.

5.2 New Projects

Early into production the partners in the business and finance departments wanted different versions of the LTV scores. For example, they wanted either bad debt or off-network charges excluded from the equation. Because of the simple, modular nature of the formula these types of requests were feasible.

The LTV project became a springboard to other projects. For instance, when an outside consultancy group prioritized the items in LTC's marketing budget based on local markets. The prioritization failed for several reasons, including: 1) it was a black box in that few knew its methodology 2) internal groups could not modify the results to produce their own analysis and 3) the prioritization only covered half of the markets. Producing a new prioritization with LTV was straightforward. Bad debt and off-network charges for example could be changed in order to

show what could happen with tighter controls and better infrastructure. It was possible to deliver special LTV analysis that only focused on new customers so LTC could see where to allocate acquisition dollars.

The only drawback of all this was that we had to do all the analysis. We only published the final results, and not all the intermediate quantities. This is something we would change if we could do the project again. Our business and finance partners would still look to us for guidance about LTV, but we could have them do the most of the work.

5.3 Regional Marketing Managers

After the project was put into production and we were educating LTC on the benefits and usage of LTV data, we ran into a substantial and justified conflict with some of the regional managers.

The issue was that LTC had expanded into areas ahead of LTC's ability to profitably support the areas: build the customer base first, and then put in the infrastructure. The regional managers in these expansion areas were drastically affected by LTV scores. This effect was on a personal level: the manager's abilities to meet their personal goals, and get their yearly bonus, were strongly effected.

We never got a full solution; the issue was still being discussed when we left the company. We were able to create partial solutions. Because of the modular nature of the LTV calculations, we were able to create an adjusted LTV that worked for these regions.

6. CONCLUSIONS

The LTV project was completed successfully. In addition to the company benefits, there were substantial career benefits: our department became known and respected as the 'LTV Department', and we gained tremendous credibility in all our other projects because of this.

Looking back, the key ingredients to the success were

- 1) The project had solid value with a rock-solid analytic base. The substance of the project really does matter. Because we had solid analytics the team believed in the project and we were able to defend the project against criticism.
- 2) The project had a high-level sponsor that was willing to go out on a limb for the project. Needing an executive sponsor is a truism that is actually true.
- 3) The team could make decisions about the project without having to go back to the sponsor. Many times (most notably with SIBU) we were negotiating with other business units, and the sponsor had to trust us to get it right.
- 4) We designed the end result to be usable.
- 5) The core design team and implementation team were the same. A hand-off between design and implementation is a natural place for projects to die.

After the LTV project, Finance initiated an LTV-like project for Activity-Based Costing. This was so that LTC could get a handle

on its costs at a very low level, which was something LTC desperately needed. However, what LTC Finance did was to first hire a consulting group for a year of design work. The consultant group had endless meetings with every group in the company, and eventually produced a massive design document. The design spec went to IT, IT replied the project would take \$16M to build, and the project was shelved.

Design teams need to understand what the implementation issues are and implementation teams need to understand what the design priorities are. If the teams do not share the same core then they need to be able to work very closely.

6.1 What We Would Do Differently

There were some things that did not go well. Topics worth further experimenting with in future projects include:

- 1) Put together a simple reporting engine running off of our desktops first. A system like this could have spotted the regional problem.
- 2) Publish all the calculational components of the LTV system, in order for users of the system to be able to customize the results.

7. ACKNOWLEDGMENTS

Our thanks go to Brendan Kitts for his preliminary review and feedback.

8. REFERENCES

- [1] D. R. Cox and D. Oakes. Analysis of Survival Data. Chapman and Hall, London, (1984).
- [2] C. Cunningham, I. Song, and P. Chen. Data warehouse design to support customer relationship management analyses. ACM CIKM 2005, DOLAP Workshop, ISBN:1-58113-977-2, (2005), 14-22
- [3] F. R. Dwyer. Customer Lifetime Valuation for Marketing Decision Making. Journal of Direct Marketing, 11, 4 (1997), pp 6-13
- [4] H. Hwang, T. Jung, and E. Suh. An LTV model and customer segmentation based on customer value: a case study on the wireless telecommunication industry. Expert Systems with Applications, 26, (2004), 181-188.
- [5] S. Rosset, and E. Neumann. Integrating Customer Value Considerations into Predictive Modeling. IEEE ICDM 2003, ISBN:0-7695-1978-4, (2003), 283.
- [6] S. Rosset, E. Neumann, U. Eick, and N. Vatnik. Customer Lifetime Value Models for Decision Support. Data Mining and Knowledge Discovery, 7, 3 (July 2003), 321-339.
- [7] S. Rosset, E. Neumann, U. Eick, N. Vatnik, and Y. Idan. Customer lifetime value modeling and its use for customer retention planning. ACM KDD 2002, (July 2002), 332-340.