XLNT consulting "turning data into dollars"

Tom's Ten Data Tips – February 2008

Data Mining for CRM

Customer Relationship Management (CRM) was an over-hyped term that has fallen from grace. Nonetheless, principles for managing the value of a portfolio of customers remain equally valid as before. Data mining can serve many roles by supporting fact based decision making to optimize customer relationships.

Throughout the customer lifecycle, data mining has been successfully applied for targeting customer acquisition, up- and/or cross-sell, customer risk management, and retention or churn management. Segmentation can enhance insight, collaborative filtering can improve the customer experience, opportunities for profitable application of data mining abound. It is for this reason that CRM has been labeled the killer application for data mining.

1. Managing The *Entire* Customer Lifecycle Works Best

Data mining can be used in various stages of the customer lifecycle. Acquisition efforts can be optimized with data mining to find an optimal subset of prospects to contact from the universe of non-customers. Cross-sell can be enhanced by finding the right customers to make an offer, or to determine the most economic channel mix. Retention efforts should be aimed at the right customers, offering promotion and/or rebates to potential churners, and *not* to those who would have stayed anyway. Again, data mining models have been successfully employed for this.

However, piecemeal optimization does *not* necessarily lead to finding a *global* best solution (*across* the customer lifecycle). Because of organizational boundaries, and departments each having their own objectives and targets, arriving at a "global best solution" can be quite difficult to achieve (see also tip #10).

2. Model Structure Should Mirror The Business Process

When considering the appropriate technique for a given business problem, it is good practice to search for a model structure that is as close as possible to the real world process. For instance, if customers

are obliged to take products in a fixed order, this sequence can be modeled with hazard rates using survival analysis.

In general, there appears to be a tendency to model processes with linear techniques (or sometimes quadratic equations). In the real world the relations between variables tend to be more intricate, where variables can often mutually influence each other. Such realities require either a path analysis or structural equation model, or alternatively a system dynamics model. Linear or quadratic models are sometimes just too simple. The tendency to opt for linear (or quadratic) models is an oversimplification that can (dangerously) distort reality. The sub prime crisis is a case in point here where interest rates, economic downturns, aggressive mortgage selling and falling house prices all interact. Einstein said: “a model should be as simple as possible, but not too simple so as to distort reality.”

If the only tool you have is a hammer, everything looks like a nail. A good data miner is versatile and eclectic. For this reason, you need a comprehensive toolkit and an assortment of technical skills.

3. Two-Stage Models: First Predict Response, Then Amounts

A classic data mining task is to predict amounts people will credit who respond to an offer. For example, a charity may aim to target high value donations, or a bank wants to attract high net worth savings customers. In such cases, experience has shown (e.g. results from the 1998 KDD Challenge cup and evaluations) that a two stage model gives the best predictive accuracy.

First you predict who is most likely to respond. Then, for responders you predict what *amount* they will contribute. This two-phase approach mimics the actual response pattern (see also tip #2). Overlaying such models is a non-obvious task because an explicit relation between model scores and ‘true’ response probabilities needs to be made. This two-phase approach may be worth the extra effort if additional predictive power has sufficient business value.

4. Predicting Leads To Efficiency Gains, Explanation Is Geared Towards Effectiveness Improvement

Data mining can be used for predicting as well as insight. Even so, a predictive model should *always* be accompanied by some form of explanation (see tip #10 last month on data preparation).

A predictive model serves to run the existing business more efficiently: its sole purpose is waste reduction. Whether it is finding an optimal subset of customers for an offer, or determining likely defaulters, or forecasting to minimize write-offs in perishable goods. An explanatory model's purpose is to innovate, if maybe only by small incremental steps. By creating refreshing insight, new avenues for creating value may emerge. Sometimes entirely new business models come from explanatory modeling. Jeff Bezos attributes his decision to launch amazon.com to an insight he had after modeling the growth of the internet in 1994.

In short, predictive models improve efficiency, doing more of the same but cheaper. Explanatory models help to rethink *what* you do, and consider new ways to exploit market opportunities and core competencies.

5. Cross-Sell Models Should Be Customer- Rather Than Product Centric

The most commonly used approach for cross sell modeling is to build a propensity model and determine the subgroup of customers for which it is profitable to make an offer. This means that you sort the customer base on their proclivity to respond, and then you determine the last customer for whom the response probability multiplied by the net present value of a sale equals the marginal cost of contact. When you run multiple propensity models in parallel, you can pick "the best" offer (see also tip #3, newsletter campaign optimization).

The practical reality of business usually calls for business rules with regard to contact strategy. Some channels might be restricted for certain customers, or a minimum interval between (potentially 'competing' or conflicting) offers is in place, etc. Often, a central calendar planning for thematic campaigns is established. *None of this is truly customer centric*. This is where customer "ownership" can beat product silos that are perforce not easily aligned (see also tip #1).

6. Riskiest Customers Are Most Profitable

It has often been shown that customers for lending products who barely pass the credit scorecard hurdle turn into the most profitable ones. The reason for this is that their credit limit utilization (CLU) tends to be highest, and interest charges weigh relatively heavy in customer profitability.

The caveat of course is that as customers become riskier, they will also generate more losses because default percentages go up. How to manage this conundrum is a largely unsolved mathematical mystery to date.

7. What Is An Optimal Channel Mix?

There has been much ado about finding an 'optimal' channel mix for different customer segments. Much of this discussion is poorly focused and misinformed. On the basis of cost/yield considerations across channels in conjunction with response functions, it is possible to calculate which customers should get an offer via email, mail, phone, or face to face. However, most of this advanced math tends to miss the true strategic considerations.

First of all, the typical mathematical solution implies that customers use one or the other channel, when most companies are actively moving towards a multi-channel approach. Customers are encouraged to examine an offer on the web to short-circuit the sales process. Sometimes customers express explicit channel preferences that a customer centric organization can hardly ignore. To simplify channel selection to a mathematical equation that optimizes contact costs does grave injustice to the strategic considerations at play.

8. Targeting Is (Currently) The Mainstay Of Data Mining For CRM

Data mining is still a young profession and although there are many more advanced applications possible, at present 'simple' targeting is where the preponderance of projects are taking place. The mining goal is to find an optimal subset within a large population. The same process takes many different forms: which customers to make an offer to, detect which transactions are likely to be fraudulent, who will default on their loans, etc.

As data mining matures, most of these applications will be 'taken for granted', and have been a part of a new generation's curriculum. Accountability for results will be a given. Most companies need to improve monitoring of their results to get there, though.

9. Data Modeling Advances Have made CRM Possible

Customer databases have been available in for some time in one form or another. However, advances in data modeling and integration in the

past 10 only have made it possible to “view” customers from an entirely new perspective.

As the marketing focus shifted from market- to customer share, a ‘new’ perspective on customer behavior was needed. What happened was that star-schema models facilitated a longitudinal perspective on the customer that was potentially available before, but extremely cumbersome to produce in older data models. It is in this sense that the way data a stored have a huge impact on what kind of phenomena are likely to be ‘seen’.

10. Customer Retention Challenges Usually Are Acquisition Failures

Data mining has been used successfully to predict customers that are likely to churn on innumerous occasions. However, being able to *predict* who is likely to churn in and of itself is not enough for success. Then you need to plan timely (!) actions to compel customers to stay rather than churn.

All too often, the challenge in coming up with effective interventions is like an uphill struggle. When a good ‘match’ between customer needs and the intrinsic product/service proposition was never there, it becomes exponentially hard to convince such customers to stay! Thus, poor (sometimes overly aggressive) acquisition strategy often leads to fruitless retention efforts. The mobile telco’s offer a painful case in point here, where the nature of the offerings make it most attractive for customers to keep switching providers after their contracts expire.

