International Journal of Advanced Computer Research (ISSN (print): 2249-7277 ISSN (online): 2277-7970)
Volume-1 Number-2 Issue-2 December 2011

A Study of Analytics Driven Solutions for Customer Targeting and Sales Force Allocation in Data Mining

Santosh Soni[1]
Assistant Professor, VITS Satna, M.P

Abstract

OnTARGET and MAP are examples of analytics-based solutions that were designed from the outset to address specific business challenges in the broad area of sales force productivity. Although they address different underlying issues, these solutions implement a common approach that is generally applicable to a broad class of operational challenges. Both solutions rely on rigorously defined data models that integrate all relevant data into a common database. Choices of the data to be included in the data model are driven both by end-user requirements as well as the need for relevant inputs to analytical models. Both business problems have a natural mapping to applications of predictive modeling: predicting the probability to purchase in the case of OnTARGET, and estimating the realistic revenue opportunity in the case of MAP. Delivering the underlying data and the analytic insights directly to frontline decision makers (sales representatives for OnTARGET and sales executives for MAP) is crucial to driving business impact, and a significant effort has been invested in developing efficient web-based tools with the necessary supporting infrastructure. In this paper we discuss several aspects and analyze them.

Keywords

OnTARGET, MAP, Data Mining, Sales Force

1. Introduction

Data mining is an ambiguous term that has been used to refer to the process of finding interesting information in large repositories of data. More precisely, the term refers to the application of special algorithms in a process built upon sound principles from numerous disciplines including statistics, artificial intelligence, machine learning, database science, and information retrieval according to Han and Kamber (2001).

Data mining algorithms are utilized in the process of pursuits variously called data mining, knowledge mining, data driven discovery, and deductive learning {Dunham (2003)}. Data mining techniques can be performed on a wide variety of data types including databases, text, spatial data, temporal data, images, and other complex data. Some areas of specialty have a name such as KDD (knowledge discovery in databases), text mining and Web mining. Most of these specialties utilize the same basic toolset and follow the same basic process and (hopefully) yield the same product – useful knowledge that was not explicitly part of the original data set.

Improving sales productivity is an essential component of driving organic growth for many major companies today. While hiring the best sales representatives is an obvious first step, it is increasingly recognized [1] that the realization of the true potential of any sales force requires that sales reps and executives be equipped with relevant IT-based tools and solutions. The past decade has seen the development of a number of customer relationship management (CRM) systems [2, 3] that provide integration and management of data relevant to the complete marketing and sales process. Sales force automation (SFA) systems [4] enable sales executives to better balance sales resources against identified sales opportunities. While it is generally (but not uniformly [5]) accepted that such tools improve the overall efficiency of the sales process, major advances in sales force productivity require not only access to relevant data, but informative, predictive analytics derived from this data. In this paper, we develop analytical approaches to address two issues relevant to sales force productivity, and describe the deployment of the resulting solutions within IBM. The first solution addresses the problem faced by sales representatives in identifying new sales opportunities at existing client accounts as well as at non-customer (“whitespace”) companies.

The analytical challenge is to develop models to predict the likelihood (or propensity) that a company will purchase an IBM product, based on analysis of previous transactions and other available third-party data. These modeling results, along with the underlying data, have been integrated in a web-based tool called on TARGET. A second, but related business challenge is to provide quantitative insight into the process of allocating sales reps to the best potential revenue-generation opportunities. In particular, we are interested in the allocation of resources to existing IBM client accounts. Here, the analytics challenge is to develop models to estimate the true revenue potential (or opportunity) at each account within IBM product groups. These models were developed as part of an internal initiative called the Market Alignment Program (MAP), in which the model-estimated revenue opportunities were validated via extensive interviews Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee with front-line sales teams.

We describe this process, and the web-based MAP tool, later in this paper. Although they address different business problems, the OnTARGET and MAP tools share a common architecture. Both employ a data model that effectively joins historical IBM transaction data with external third-party data, thereby presenting a holistic view of each client in terms of their past history with IBM as well as their external "firmographic" information like sales, number of employees, and so on. Both systems exploit this linked data to build the models described above and in the sections below. Given the different business objectives, the tools employ different web-based user interfaces; however, both interfaces are designed to facilitate easy navigation and location of the relevant analytical insights and underlying data. In the following section, we describe the OnTARGET project, its data model, and overall system design motivated by the business requirements. We then describe the propensity models at the heart of OnTARGET. Turning to the MAP project in Section 4, we discuss the MAP business process, and describe the differences in the MAP tool design relative to OnTARGET, again motivated by the different business objectives of MAP. Section 5 describes the MAP revenue-opportunity models. Finally, we describe the deployment of these systems, and discuss the operational impact against their respective business objectives.

We provide here an overview of executing data mining services . The rest of this paper is arranged as follows: Section 2 introduces Advent of data mining; Section 3 describes about working of data mining; Section 4 shows the technique used; Section 5 describes the A Sales Force Allocation Tool. Section 6 describes Conclusion and outlook.

2. Advent of Data Mining

Data mining derives its name from the similarities between searching for valuable business information in a large database for example, finding linked products in gigabytes of store scanner data and mining a mountain for a vein of valuable ore. Both processes require either sifting through an immense amount of material, or intelligently probing it to find exactly where the value resides. Given databases of sufficient size and quality, data mining technology can generate new business opportunities by providing these capabilities:

- Automated prediction of trends and behaviors. Data mining automates the process of finding predictive information in large databases. Questions that traditionally required extensive hands-on analysis can now be answered directly from the data — quickly. A typical example of a predictive problem is targeted marketing. Data mining uses data on past promotional mailings to identify the targets most likely to maximize return on investment in future mailings. Other predictive problems include forecasting bankruptcy and other forms of default, and identifying segments of a population likely to respond similarly to given events.
- Automated discovery of previously unknown patterns. Data mining tools sweep through databases and identify previously hidden patterns in one step. An example of pattern discovery is the analysis of retail sales data to identify seemingly unrelated products that are often purchased together. Other pattern discovery problems include detecting fraudulent credit card transactions and identifying anomalous data that could represent data entry keying errors.

Data mining techniques can yield the benefits of automation on existing software and hardware platforms, and can be implemented on new systems, as existing platforms are upgraded and new products developed. When data mining tools are implemented on high performance parallel processing systems, they can analyze massive databases in minutes. Faster processing means that users can automatically experiment with more models to understand complex data. High speed makes it practical for users to analyze huge quantities of data. Larger databases, in turn, yield improved predictions.

Databases can be larger in both depth and breadth:

- More columns. Analysts must often limit the number of variables they examine when doing hands-on analysis due to time constraints. Yet variables that are discarded because they seem unimportant may carry information about unknown patterns. High performance data mining allows users to explore the full depth of a database, without pre selecting a subset of variables.

- More rows. Larger samples yield lower estimation errors and variance, and allow users to make inferences about small but important segments of a population.

A recent Gartner Group Advanced Technology Research Note listed data mining and artificial intelligence at the top of the five key technology areas that "will clearly have a major impact across a wide range of industries within the next 3 to 5 years."2 Gartner also listed parallel architectures and data mining as two of the top 10 new technologies in which companies will invest during the next 5 years. According to a recent Gartner HPC Research Note, "With the rapid advance in data capture, transmission and storage, large-systems users will increasingly need to implement new and innovative ways to mine the after-market value of their vast stores of detail data, employing MPP [massively parallel processing] systems to create new sources of business advantage (0.9 probability).

The most commonly used techniques in data mining are:

- Artificial neural networks: Non-linear predictive models that learn through training and resemble biological neural networks in structure.
- Decision trees: Tree-shaped structures that represent sets of decisions. These decisions generate rules for the classification of a dataset. Specific decision tree methods include Classification and Regression Trees (CART) and Chi Square Automatic Interaction Detection (CHAID).
- Genetic algorithms: Optimization techniques that use process such as genetic combination, mutation, and natural selection in a design based on the concepts of evolution.
- Nearest neighbor method: A technique that classifies each record in a dataset based on a combination of the classes of the k record(s) most similar to it in a historical dataset (where k ³ 1). Sometimes called the k-nearest neighbor technique.

Rule induction: The extraction of useful if-then rules from data based on statistical significance.
Many of these technologies have been in use for more than a decade in specialized analysis tools that work with relatively small volumes of data. These capabilities are now evolving to integrate directly with industry-standard data warehouse and OLAP platforms. The appendix to this white paper provides a glossary of data mining terms.

3. Working of Data Mining

How exactly is data mining able to tell you important things that you didn't know or what is going to happen next? The technique that is used to perform these feats in data mining is called modeling. Modeling is simply the act of building a model in one situation where you know the answer and then applying it to another situation that you don't. For instance, if you were looking for a sunken Spanish galleon on the high seas the first thing you might do is to research the times when Spanish treasure had been found by others in the past. You might note that these ships often tend to be found off the coast of Bermuda and that there are certain characteristics to the ocean currents, and certain routes that have likely been taken by the ship's captains in that era. You note these similarities and build a model that includes the characteristics that are common to the locations of these sunken treasures. With these models in hand you sail off looking for treasure where your model indicates it most likely might be given a similar situation in the past. Hopefully, if you've got a good model, you find your treasure.

This act of model building is thus something that people have been doing for a long time, certainly before the advent of computers or data mining technology. What happens on computers, however, is not much different than the way people build models. Computers are loaded up with lots of information about a variety of situations where an answer is known and then the data mining software on the computer must run through that data and distill the characteristics of the data that should go into the model. Once the model is built it can then be used in similar situations where you don't know the answer. For example, say that you are the director of marketing for a telecommunications company and you'd like to acquire some new long distance phone customers. You could just randomly go out and mail coupons to the general population - just as you could randomly sail the seas looking for sunken treasure. In neither case would you achieve the results you desired and of course you have the opportunity to do much better than random - you could use your business experience stored in your database to build a model.

The architecture may have following major component which is shown in Fig 1.

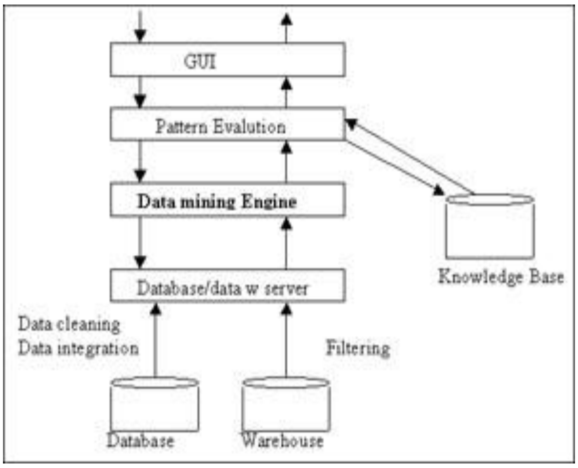

Fig 1: Data Mining Architecture

4. Techniques of Data Mining

This overview provides a description of some of the most common data mining algorithms in use today. We have broken the discussion into two sections, each with a specific theme:

- Classical Techniques: Statistics, Neighborhoods and Clustering
- Next Generation Techniques: Trees, Networks and Rules

Each section will describe a number of data mining algorithms at a high level, focusing on the "big picture" so that the reader will be able to understand how each algorithm fits into the landscape of data mining techniques. Overall, six broad classes of data mining algorithms are covered. Although there are a number of other algorithms and many variations of the techniques described, one of the algorithms from this group of six is almost always used in real world deployments of data mining systems. After the collapse of the Internet bubble, corporate growth has returned to what are likely to be normal growth rates for the next few years. Since the broad market is likely to grow in aggregate at rates only slightly higher than GDP, companies will need to generate organic revenue growth at rates greater than the market overall to remain competitive. One approach is to pursue growth opportunities in emerging markets. But it is also necessary to generate significant growth in a company's core businesses and markets. This requires a renewed focus on identifying and closing new sales opportunities with existing clients, as well as finding new companies that will be receptive to the company's core offerings. Improving sales force productivity is essential to both objectives. Early in the OnTARGET project, we spoke to a number of leading sales professionals and sales leaders about potential It enabled tools that they believed could enhance sales productivity. One common sentiment is that sales people are often forced to use multiple tools and processes that not only fail to provide the relevant information needed to do their jobs better, but also take valuable time away from actual sales activities. While some cross-sell models were available for use by the sales teams, these analytics were often delivered via spreadsheets and lacked integration with important underlying data needed to understand the full client sales history and requirements.

I. References a large universe of existing clients and potential new clients,
II. Incorporates relevant data that may require multiple existing tools to access,
III. Includes analytical models to help identify the best sales opportunities, and
IV. Integrates all such data for each company under a single user interface designed by end users to facilitate easy navigation.

As discussed further in Section, OnTARGET is now used by 7,000 IBM sales representatives, and has largely replaced the previous set disparate sales tools and spreadsheets. The success of OnTARGET is due in large measure to our ability to deliver these key capabilities directly to the front-line sales force. In the rest of this section, we discuss specific design decisions and implementations in light of these requirements. In particular, we discuss the types of data selected for inclusion in the tool, the integration of this data in the overall OnTARGET system, and the design of the user interface.

5. A Sales Force Allocation Tool

In this section, we discuss the second initiative, the Market Alignment Program (MAP), mentioned in the Introduction. We first describe the project objectives, and then delineate the business process we developed to address these objectives. An integral part of the MAP process is the validation of analytical estimates via an extensive set of workshops conducted with sales leaders. These interviews rely heavily on a web-based tool to convey the relevant information, as well as to capture the expert feedback on the analytical models. We describe the design of the MAP tool infrastructure, its data model, and the key characteristics of the user interface. In 2010, vendors turned their attention to meeting the needs of power users after ten years of enhancing reporting and dashboard solutions for casual users. As a result, the number of analytical tools on the market has exploded. Analytical tools come in all shapes and sizes. Analysts generally need one of every type of tool. Just as you wouldn't hire a carpenter to build an addition to your house with just one tool, you don't want to restrict an analyst to just one analytical tool. Like a carpenter, an analyst needs a different tool for every type of job they do. For instance, a typical analyst might need the following tools:

Excel to extract data from various sources, including local files, create reports, and share them with others via a corporate portal or server (managed Excel). BI Search tools to issue ad hoc queries against a BI tool's metadata.

Planning tools (including Excel) to create strategic and tactical plans, each containing multiple scenarios.

Mash boards and ad hoc reporting tools to create ad hoc dashboards and reports on behalf of departmental colleagues Visual discovery tools to explore data in one or more sources of data and create interactive dashboards on behalf of departmental colleagues. Multidimensional OLAP (MOLAP) tools to explore small and medium sets of data dimensionally at the speed of thought and run complex dimensional calculations. Relational OLAP tools to explore large sets of data dimensionally and run complex calculations Text analytics tools to parse text data and put it in a relational structure for analysis.

Data mining tools to create descriptive and predictive models. Hadoop and MapReduce to process large volumes of unstructured and semi-structured data in a parallel environment.

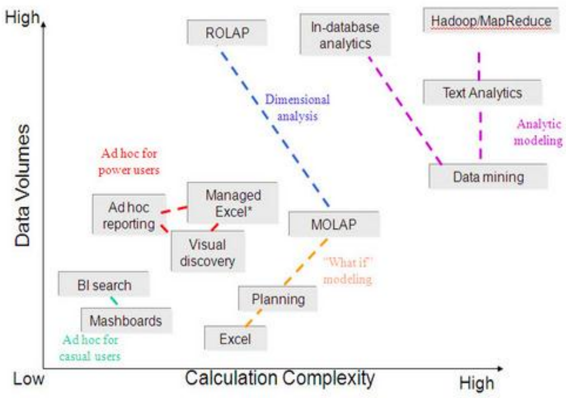

Fig 2: Types of Analytical Tool

6. Conclusion and Outlook

The solutions have been deployed across multiple geographic regions, with a strong focus on capturing and quantifying the business impact of the initiatives. Indeed, we have field evidence that the analytical models developed for OnTARGET are predictive. MAP is a more recent initiative, but preliminary evidence suggests that sales force allocations made within the MAP process are leading to measurable improvements in sales efficiency. Finally, although we have implemented these solutions within IBM, we believe that the underlying methodologies, business processes, and potential impact are relevant to enterprise sales organizations in many other global industries.

References

[1] Berry, M.J.A. and G.S. Linoff, Data Mining Techniques: For Marketing, Sales, and Customer Relationship Management.2004: John Wiley.

[2] Breiman, L Classification and Regression Trees.1984, Belmont, CA: Wadsworth International Group.

[3] Hastie, T., R. Tibshirani, and J.H. Friedman, Elements of Statistical Learning,. 2003: Springer.

[4] IBM Rational Data Architect, http://www.ibm.com/software.

[5] IBM Web Sphere DataStage,http://www.ibm.com/software/data/integration/datastage/.Koenker, R., Quantile Regression. Econometric Society Monograph Series. 2005, Cambridge: Cambridge University Press.