

A Study on Scientific application of Data Mining

H S Venkatesh Prasad¹, Madhu B K², Lokesha V³

¹Lecturer in Computer Science, FMKMC College, Madikeri. Coorg, Karnataka

²Professor, Dept of CSE, GSSIT, Bangalore

³Director, Dept of MCA, ACIT, Bangalore

Email: hsvp_tvk@yahoo.co.in

Abstract

The relatively new discipline of data mining is most often applied to extraction of useful knowledge from business data. However, it is also useful in some scientific applications where this more empirical approach complements traditional data analysis. The example of machine learning from air quality data illustrates this alternative.

1. Introduction

Data Mining is the essential ingredient in the more general process of Knowledge Discovery in Databases (KDD). The idea is that by automatically sifting through large quantities of data it should be possible to extract nuggets of knowledge. Data mining has become fashionable, not just in computer science (journals & conferences), but particularly in business IT. (An example is its promotion by television advertising [1].) The emergence is due to the growth in data warehouses and the realisation that this mass of operational data has the potential to be exploited as an extension of *Business Intelligence*. Data Mining is a multidisciplinary field, drawing work from areas including database management systems, artificial intelligence, machine learning, neural networks, statistics, pattern recognition, knowledge-based systems, knowledge acquisition, information retrieval, high-performance computing, and data visualization. There are a number of definitions of Data Mining in the literature. However, they all have in common the following; "extraction", "knowledge", and "large data." (Abbass, 2005). Data mining refers to the extracting or "mining" knowledge from large amount of data. The process of performing data analysis may uncover important data patterns, contributing greatly to business strategies, knowledge bases, and scientific and medical research. The exploration and analysis, by automatic or semiautomatic means, of large quantities of data in order to discover meaningful patterns and rules.

The rules include the iterative process of detecting and extracting patterns from large databases. It lets us identify "signatures" hidden in large databases, as well as learn from repeated examples. The extraction of implicit, previously unknown, and potentially useful information from data is the ultimate goal of any statically viable approach. The idea is to build computer programs that sift through databases automatically, seeking regularities or patterns. Strong

patterns, if found, will likely generalize to make accurate predictions on future data. Data Mining is also the search for valuable information in large volumes of data. It is a cooperative effort of humans and computers. Humans design databases, describe problems and set goals.

2. Why is Data Mining Different?

Data mining is more than just conventional data analysis. It uses traditional analysis tools (like statistics and graphics) plus those associated with artificial intelligence (such as rule induction and neural nets). It is all of these, but different. It is a distinctive approach or attitude to data analysis. The emphasis is not so much on extracting facts, but on generating hypotheses. The aim is more to yield questions rather than answers. Insights gained by data mining can then be verified by conventional analysis.

3. The Data Management Context

"Information Technology" was originally "Data Processing". Computing in the past gave prominence to the processing algorithms - data were subservient. Typically, a program processed input data tapes (such as master and detail records) in batch to output a new data tape that incorporated the transactions. The structure of the data on the tapes reflected the requirements of the specific algorithm. It was the era of Jackson Structured Programming. The concept of database broke away from this algorithm-centric view. Data assumed an existence independent of any programs. The data could be structured to reflect semantics of relationships in the real world. One had successively hierarchical, network, relational and object data models in commercial database management systems, each motivated by the desire to model better the structure of actual entities and their relationships. A database is extensional, storing many facts. Some information is intensional; that is, it manifests as rules. Some limited success was achieved with deductive databases that stored and manipulated rules, as for example in Prolog based systems. This encouraged *Expert Systems*. However, it was hard to achieve solid success. A main difficulty was the knowledge elicitation bottleneck: how to convert the thought processes of domain experts into formal rules in a computer.

Data mining offers a solution: automatic rule extraction. By searching through large amounts of data, one hopes to find sufficient instances of an

association between data value occurrences to suggest a statistically significant rule. However, a domain expert is still needed to guide and evaluate the process and to apply the results.

4. Business Data Analysis

Popular commercial applications of data mining technology are, for example, in direct mail targeting, credit scoring, churn prediction, stock trading, fraud detection, and customer segmentation. It is closely allied to data warehousing in which large (gigabytes) corporate databases are constructed for decision support applications. Rather than relational databases with SQL, these are often multi-dimensional structures used for so called *on-line analytical processing* (OLAP). Data mining is a step further from the directed questioning and reporting of OLAP in that the relevant results cannot be specified in advance.

5. Scientific Data Analysis

Rules generated by data mining are empirical - they are not physical laws. In most research in the sciences, one compares recorded data with a theory that is founded on an analytic expression of physical laws. The success or otherwise of the comparison is a test of the hypothesis of how nature works expressed as a mathematical formula. This might be something fundamental like an inverse square law. Alternatively, fitting a mathematical model to the data might determine physical parameters (such as a refractive index). On the other hand, where there are no general theories, data mining techniques are valuable, especially where one has large quantities of data containing noisy patterns. This approach hopes to obtain a theoretical generalisation automatically from the data by means of induction, deriving empirical models and learning from examples. The resultant theory, while maybe not fundamental, can yield a good understanding of the physical process and can have great practical utility.

6. Scientific Applications

In a growing number of domains, the empirical or black box approach of data mining is good science. Three typical examples are:

6.1. Sequence analysis in bioinformatics. Genetic data such as the nucleotide sequences in genomic DNA are digital. However, experimental data are inherently noisy, making the search for patterns and the matching of sub-sequences difficult. Machine learning algorithms such as artificial neural nets and hidden Markov chains are a very attractive way to tackle this computationally demanding problem.

6.2. Classification of astronomical objects. The thousands of photographic plates that comprise a large survey of the night sky contain around a billion faint objects. Having measured the attributes of each object, the problem is to classify each object as a

particular type of star or galaxy. Given the number of features to consider, as well as the huge number of objects, decision-tree learning algorithms have been found accurate and reliable for this task.

6.3. Medical decision support. Patient records collected for diagnosis and prognosis include symptoms, bodily measurements and laboratory test results. Machine learning methods have been applied to a variety of medical domains to improve decision-making. Examples are the induction of rules for early diagnosis of rheumatic diseases and neural nets to recognise the clustered micro-calcifications in digitised mammograms that can lead to cancer. The common technique is the use of data instances or cases to generate an empirical algorithm that makes sense to the scientist and that can be put to practical use for recognition or prediction.

7. Example of Predicting Air Quality

To illustrate the data mining approach, both advantages and disadvantages, this section describes its application to a prediction of metropolitan air pollution.

7.1 Motivation. One needs an understanding of the behaviour of air pollution in order to predict it and then to guide any action to ameliorate it. Calculations with dynamical models are based on the relevant physics and chemistry. An interesting research and development project pursuing this approach is DECAIR [5]. This concerns a generic system for exploiting existing urban air quality models by incorporating land use and cloud cover data from remote sensing satellite images.

To help with the design and validation of such models, a complementary approach is described here. It examines data on air quality empirically. Data mining and, in particular, machine learning techniques are employed with two main objectives:

7.1.1. To improve our understanding of the relevant factors and their relationships, including the possible discovery of non-obvious features in the data that may suggest better formulations of the physical models;

7.1.2. To induce models solely from the data so that dynamical simulations might be compared to them and that they may also have utility, offering (short term) predictive power.

7.2 Data Preparation

Before trying to apply machine learning and constructing a model, there are three quite important stages of data preparation. The data need to be *cleaned*, *explored* and *transformed*. In typical applications, this can be most of the overall effort involved.

a) Cleaning

Though not elaborated here, but commonly a major part of the KDD process is data cleaning. In this case, one is concerned with imposing consistent formats for dates and times, allowing for missing data, finding duplicated data, and weeding out bad data – the latter are not always obvious. The treatment of missing or

erroneous data needs application dependent judgement.

b) Exploration

Another major preliminary stage is a thorough examination of the data to acquire familiarity and understanding. One starts with basic statistics - means, distributions, ranges, etc - aiming to acquire a feeling for data quality. Other techniques such as sorting, database queries and especially exploratory graphics help one gain confidence with the data.

c) Transformation

The third preparatory step is dataset sampling, summarisation, transformation and simplification. Working with only a sample of the full data, or applying a level of aggregation, may well yield insights and results that are quicker (if not even discernible at all) than with the complete data source. In addition, transforming the data by defining new variables to work with can be a crucial step. Thus one might, for instance, calculate ratios of observations, normalise them, or partition them into bins, bands or classes.

8. Conclusions

Work so far supports the common experience in data mining that most of the effort is in data preparation and exploration. The data must be cleaned to allow for missing and bad measurements. Detailed examination leads to transforming the data into more effective forms. The modelling process is very iterative, using statistics and visualisation to guide strategy. The temporal dimension with its lagged correlations adds significantly to the search space for the most relevant parameters. Investigation that is more extensive is needed to establish under what circumstances data mining might be as effective as dynamical modelling. (For instance, urban air quality varies greatly from street to street depending on buildings and traffic.) A feature of data mining is that it can “short circuit” the post-interpretation of the output of numerical simulations by directly predicting the probability of exceeding pollution thresholds. A drawback is the need for large datasets in order to provide enough high pollution episodes for reliable rule induction. More generally, data mining analysis is useful to provide a reference model in the validation of physically based simulation calculations. Is data mining as useful in science as in commerce. Certainly, data mining in science has much in common with that for business data. One difference, though, is that there is a lot of existing scientific theory and knowledge. Hence, there is less chance of knowledge emerging purely from data. However, empirical results **can** be valuable in science (especially where it borders on engineering) as in suggesting causality relationships or for modelling complex phenomena. Another difference is that in commerce, rules are soft - sociological or cultural – and assume consistent behaviour. For example, the plausible myth that “30% of people who buy babies’

nappies also buy beer” is hardly fundamental, but one might profitably apply it as a selling tactic (until perhaps the fashion changes from beer to lager). On the other hand, scientific rules or laws are, in principle, testable objectively. Any results from data mining techniques must sit within the existing domain knowledge.

8. References

- [1] Abbass, Hussein A. , Ruhul A. Sarker, Charles Sinclair Newton Data Mining: A Heuristic Approach, Idea group Publishing inc. , 2005Daimi, Kevin, Data Mining software, University of Detroit Mercy, 2007
- [2] Grossman, Robert L. ,DataMining for Scientific and engineering applications, Springer Publisher, 2005
- [3] Han, Jaiwei, Micheline Kamber, Data Mining: concepts and techniques, Elsevier publishing, 2003Squire, Linda, What is Data Mining, SBSS-DAMA-NCR
- [4] M. Ankerst, C. Elsen, M. Ester, and H.-P. Kriegel. Visual classification: An interactive approach to decision tree construction. In *Proc. 1999 Int. Conf. Knowledge Discovery and Data Mining (KDD'99)*, pages 392{396, San Diego, CA, Aug. 1999.
- [5] C. C. Aggarwal. *Data Streams: Models and Algorithms*. Kluwer Academic, 2006.
- [6] C. C. Aggarwal, J. Han, J. Wang, and P. S. Yu. A framework for clustering evolving data streams. In *Proc. 2003 Int. Conf. Very Large Data Bases (VLDB'03)*, pages 81{92, Berlin, Germany, Sept. 2003.
- [7] C. C. Aggarwal, J. Han, J. Wang, and P. S. Yu. A framework for projected clustering of high dimensional data streams. In *Proc. 2004 Int. Conf. Very Large Data Bases (VLDB'04)*, pages 852{863, Toronto, Canada, Aug. 2004.
- [8] James Allan. *Topic Detection and Tracking: Event-Based Information Organization*. Kluwer Academic Publishers, Norwell, MA, USA, 2002.
- [9] B. Babcock, S. Babu, M. Datar, R. Motwani, and J. Widom. Models and issues in data stream systems. In *Proc. 2002 ACM Symp. Principles of Database Systems (PODS'02)*, pages 1{16, Madison, WI, June 2002.
- [10] M. W. Berry. *Survey of Text Mining: Clustering, Classification, and Retrieval*. Springer, 2003.
- [11] I. Bhattacharya and L. Getoor. A latent dirichlet model for unsupervised entity resolution. In *Proc.2006 SIAM Int. Conf. Data Mining (SDM'06)*, Bethesda, MD, April 2006.
- [12] P. Bajcsy, J. Han, L. Liu, and J. Yang. Survey of bio-data analysis from data mining perspective. In Jason T. L. Wang, Mohammed J. Zaki, Hannu T. T. Toivonen, and Dennis Shasha, editors, *DataMining in Bioinformatics*, pages 9{39. Springer Verlag, 2004.
- [13] B.-C. Chen, L. Chen, Y. Lin, and R. Ramakrishnan. Prediction cubes. In *Proc. 2005 Int. Conf. VeryLarge Data Bases (VLDB'05)*, pages 982{993, Trondheim, Norway, Aug. 2005.
- [14] M. Craven, D. DiPasquo, D. Freitag, A. McCallum, T. M. Mitchell, K. Nigam, and S. Slattery. Learning to construct knowledge bases from the world wide web. *Artificial Intelligence*.