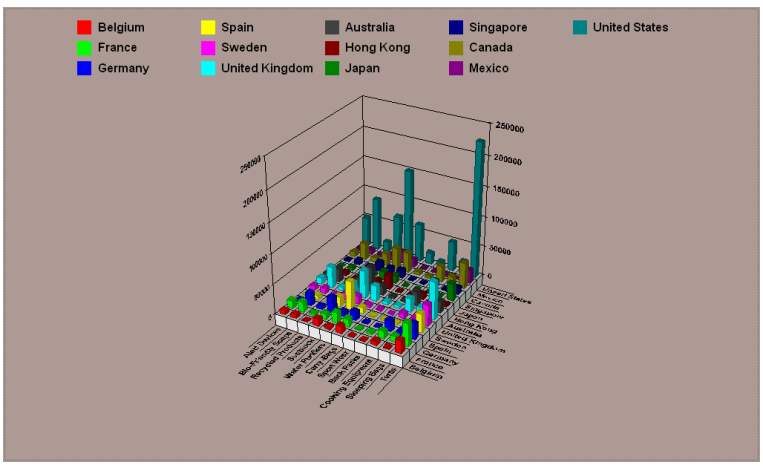

Figure 3 Graphical output of the Characterizer of DBMiner

generalized cuboid in which each attribute contains no more than but close to the threshold value of the corresponding attribute. Then the concepts in the *contrasting* class(es) are generalized to the same level as those in the prime target cube, forming the *prime contrasting cube*. Finally, the information in these two classes is used to generate qualitative or quantitative discriminant rules.

Moreover, interactive drill-down and roll-up can be performed synchronously in both target class and contrasting class(es) in a similar way as in characterization.

How can multi-level characterization and comparison be integrated into OLAP mining? Since at each step of drill-down or roll-up, characterization and comparison produce a new cuboid, with the same data structure, it is inherently suitable for integration with OLAP mining. That is, any mining module can treat the result of characterization and comparison as a data cube and mining can be performed directly on such a resulting cube. Furthermore, for any mining result on which cubing operations can be performed, characterization and comparison can be done as well.

4.2 OLAP-based association

Based on many studies on efficient mining of association rules (Agrawal & Srikant 1994, Srikant & Agrawal 1995, Han & Fu 1995), a multiple-level association rule miner (called "associator") has been implemented in DBMiner. An output of the associator is shown in Figure 4.

Different from mining association rules in transaction databases, a relational

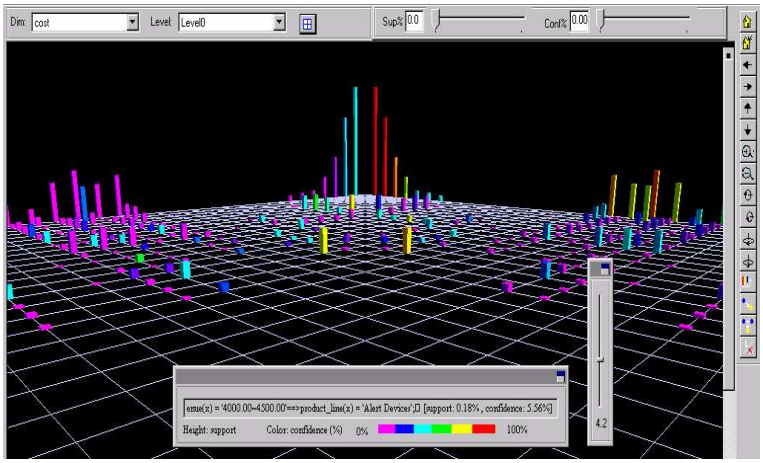

Figure 4 Graphical output of the Associator of DBMiner

associator may find two kinds of associations: *inter-attribute association* and *intra-attribute association*. The former is an association among different attributes; whereas the latter is an association within one or a set of attributes formed by grouping of another set of attributes. This is illustrated in the following example.

Example 2. Suppose the "course_taken" relation in a university database has the following schema:

$$course_taken = (student_id, course, semester, grade).$$

Intra-attribute association is the association among one or a set of attributes formed by grouping another set of attributes in a relation. For example, the associations between each student and his/her course performance is an intra-attribute association because a set of attributes, "$course, semester, grade$", are grouped according to $student_id$, for mining associations among the courses taken by each student. From a relational database point of view, a relation so formed is a nested relation obtained by nesting "$(course, semester, grade)$" with the same $student_id$. Therefore, an intra-attribute association is an association among the nested items in a nested relation.

Inter-attribute association is the association among a set of attributes in a *flat* relation. For example, the association between *course* and *grade*, such as "*the courses in computing science tend to give good grades*", is an inter-attribute association.

Two associations require different mining algorithms.

For mining intra-attribute associations, a data relation can be transformed

into a nested relation in which the tuples which share the same values in the nesting attributes are merged into one. For example, the *course_taken* relation can be folded into a nested relation with the schema,

$$course_taken = (student_id, course_history)$$
$$course_history = (course, semester, grade).$$

With such transformation, it is easy to derive association rules like "*90% senior CS students tend to take at least three CS courses at 300-level or up in each semester*". Since the nested tuples (or values) can be viewed as data items in the same transaction, the methods for mining association rules in transaction databases, such as Apriori (Agrawal & Srikant 1994), can be applied to such transformed relations in relational databases. □

Moreover, it is preferable to have some user-specified constraints to guide an association rule mining process. Such constraints can be specified in a *meta-rule* (or *meta-pattern*) form (Kamber, Han & Chiang 1997), which confines the search to specific forms of rules. For example, a meta-rule "$P(x, y) \rightarrow Q(x, y, z)$", where P and Q are predicate variables matching different properties in a database, can be used as a rule-form constraint in the search.

The multi-dimensional data cube structure facilitates efficient mining of multi-level, inter-attribute association rules. A count cell of a cube stores the number of occurrences of the corresponding multi-dimensional data values; whereas a dimension count cell stores the sum of counts of the whole dimension. With this structure, it is straightforward to calculate the measures such as *support* and *confidence* of association rules based on the values in these summary cells. A set of such cuboids, ranging from the least generalized one to rather high level ones, facilitate mining of association rules at multiple levels of abstraction.

Association mining module generates a set of association rules. Since the rule form is quite different from a data cube structure, it is not easy to integrate the mining results with other mining/cubing processes. One choice is to take any rule which contains a few connected nodes as a cuboid, from which characteristics can be displayed and drill-down or roll-up can be performed. Another choice is to take a node as a one-dimensional cuboid and show data distribution and add additional attributes for cubing and mining. The third choice is to simply backtrack to a point where cubing and other mining operations can take place naturally.

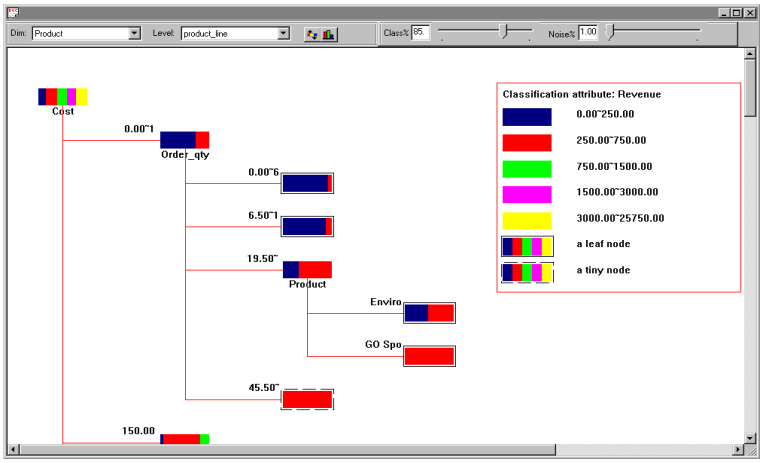

Figure 5 Graphical output of the Classifier of DBMiner

4.3 OLAP-based classification

Data classification is to develop a description or model for each class in a database, based on the features present in a set of class-labeled training data.

There have been many classification methods studied, including decision-tree methods, such as ID-3 and C4.5 (Quinlan 1993), statistical methods, neural networks, rough sets, etc. Recently, some database-oriented classification methods have also been investigated (Mehta, Agrawal & Rissanen 1996).

Our classification method consists of four steps: (1) collection of the relevant set of data and partitioning of the data into training and test data, (2) analysis of the relevance of the attributes, (3) construction of classification (decision) tree, and (4) test of the effectiveness of the classification using the test data set.

Attribute relevance analysis is performed based on the analysis of an uncertainty measurement, a measurement which determines how much an attribute is in relevance to the class attribute. Other measurements, such as entroy-based information gain (Quinlan 1993) and Gini index (Mehta et al. 1996), can be used for relevance analysis as well. Several top-most relevant attributes are retained for classification analysis; whereas the weakly or irrelevant attributes are not considered in the subsequent classification process.

In the classification process, our classifier adopts a generalization-based decision-tree induction method which integrates OLAP data cube technology with a decision-tree induction technique, by first performing minimal generalization on the set of training data, and then performing decision tree induction on the generalized data.

Since a generalized cell comes from the generalization of a number of orig-

inal cells, the *count* information is associated with each generalized cell and plays an important role in classification. To handle noise and exceptional data and facilitate statistical analysis, two thresholds, *classification threshold* and *exception threshold*, are introduced. The former is used for justification whether it is needed to continue classification on a node if a significant set of the examples of the node belongs to a single class; whereas the latter is used to terminate further classification on a node if the node contains only a negligible number of examples.

There are several alternatives for doing generalization before classification: A data set can be generalized to either a minimally generalized abstraction level, an intermediate abstraction level, or a rather high abstraction level. Too low an abstraction level may result in scattered classes, bushy classification trees, and difficulty at concise semantic interpretation; whereas too high a level may result in the loss of classification accuracy.

For OLAP mining, classification can be associated with other cubing and mining functions as follows. For any cubing result, one attribute can be selected as class attribute and classification can be performed on the corresponding cuboid in the same way as our cube-based classification process. For any classification result, each class node can be treated as a portion of the cube selected by the class constraint. Subsequent cubing and mining operations can be performed on the selected class.

The multi-level classification process has been implemented in the DBMiner system. An output of the DBMiner classifier is shown in Figure 5.

4.4 OLAP-based prediction

A predictor predicts data values or value distributions on the attributes of interest based on similar groups of data in the database. For example, one may predict the amount of research grants that an applicant may receive based on the data about the similar groups of researchers.

The power of data prediction should be confined to the ranges of numerical data or the nominal data generalizable to only a small number of categories. It is unlikely to give reasonable prediction on one's name or social insurance number based on other persons' data.

For successful prediction, the factors (or attributes) which strongly influence the values of the attributes of interest should be identified first. This can be done by the analysis of data relevance or correlations by statistical methods, decision-tree classification techniques, or be simply based on expert judgement. To analyze attribute relevance, the uncertainty measurement similar to the method used in our classifier is applied. This process ranks the

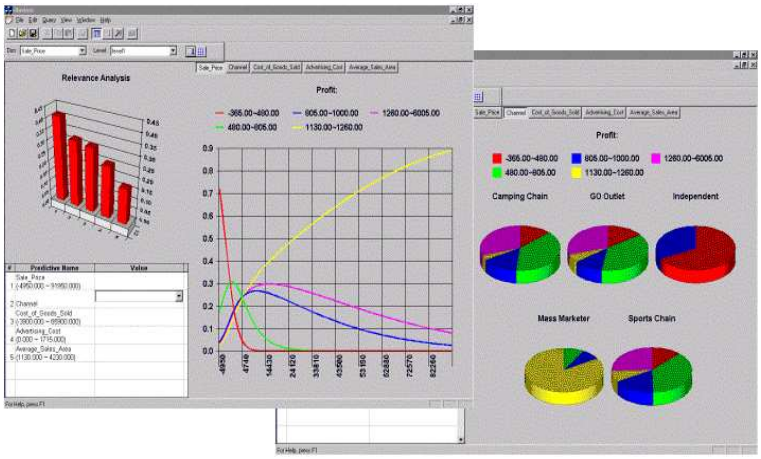

Figure 6 Graphical output of the Predictor of DBMiner: numeric predictive attribute (left) and categorical predictive attribute (right)

relevance of all the attributes selected and only the highly ranked attributes will be used in the prediction process.

After the selection of highly relevant attributes, a generalized linear model has been constructed which can be used to predict the value or value distribution of the predicted attribute. If the predictive attribute is a numerical data, a set of curves are generated, each indicating the trend of likely changes of the value distribution of the predicted attribute. If the predictive attribute is a categorical data, a set of pie charts are generated, each indicating the distributions of the value ranges of the predicted attribute.

When a query probe is submitted, the corresponding value distribution of the predicted attribute can be plotted based on the curves or pie charts generated above. Therefore, the values in the set of highly relevant predictive attributes can be used for trustable prediction.

The prediction output has two forms of presentation: curve graph and pie chart depending whether the predictive attribute is a numeric attribute or a categorical attribute. When the predictive attribute is a numeric one, the output is a set of curves as shown in the left half of Figure 6; whereas when the predictive attribute is a categorical one, the output is a set of pie charts as shown in the right half of Figure 6.

OLAP mining can be integrated with the prediction module as follows. For any predicted class, the class can be identified by a class selection criteria and its characteristics can be displayed. Then cubing operations can be performed on such a selected cuboid. Alternatively, one may backtrack to a point before prediction is performed and continue the exploration of other features of the previously selected data cube.

4.5 OLAP-based clustering analysis

Data clustering, also viewed as "unsupervised learning", is a process of partitioning a set of data into a set of classes, called *clusters*, with the members of each cluster sharing some interesting common properties. A good clustering method will produce high quality clusters, in which the intra-class (i.e., intra-cluster) similarity is high and inter-class similarity is low.

Clustering has many interesting applications. For example, it can be used to help marketers discover distinct groups in their customer bases and develop targeted marketing programs.

Data clustering has been studied in statistics, machine learning and data mining with different methods and emphases. Many clustering methods have been developed and applied to various domains, such as data classification and image processing.

Data mining applications deal with large high dimensional data, and frequently involve categorical domains with concept hierarchies. However, most of the existing data clustering methods can only handle numeric data, or cannot produce good quality results in the case where categorical domains are present.

Our cluster analyzer is based on the well-known *k-means* paradigm. Comparing to the other clustering methods, the *k-means* based methods are promising for their efficiency in processing large data sets. However, their use is often limited to numeric data. To adequately reflect categorical domains, we have developed a method of encoding concept hierarchies. This enables us to define a dissimilarity measure that not only takes into account both numeric and categorical attributes, but also at multiple levels. Due to these modifications, our cluster analyzer can cluster large data sets with mixed numeric and categorical attributes in a way similar to *k-means*. It can also perform multi-level clustering and select a desired level by a comparison of the clustering quality at different levels. On the other hand, the user or the analyst can direct the clustering process by either selecting a set of relevant attributes for the requested clustering query, or assigning a weight factor to each attribute, or both, so that increasing the weight of an attribute increases the likelihood that the algorithm will cluster according to that attribute.

OLAP mining can be integrated with cluster analyzer as follows. For any cluster so obtained, its characteristics can be displayed and cubing/mining operations can be performed on such a selected cluster. Alternatively, one may backtrack to a point before clustering is performed and continue the exploration of other features of the previously selected data cube.

4.6 Backtracking and comparative mining analysis

Backtracking is convenient for OLAP mining since a user may like to tentatively dig deep following some mining paths and later try alternatives if there have not been desired interesting patterns found.

One suggested technique for implementation of backtracking in OLAP mining is as follows. First, a status vector should be saved in a backtrack stack (if the backtrack pattern is simply tracing back step by step) or a backtrack list (if position marking or other traversal patterns is desired). The cuboid(s) associated with the vector should also be saved on the disk and linked with the vector. When backtracking, such a stack/list is used to trace back the appropriate status pointers. When a session is complete, all the saved backtrack pointers, and their associated vectors and cuboids should be deleted to release the disk space occupied.

Comparative mining analysis can be implemented similarly. For comparative analysis of two mining tasks, one needs to fork a new path of mining/cubing process stream and show both streams in separate windows. Comparative displays can also be synchronized when necessary by bundling together the two similar mining/cubing processes.

5 DISCUSSION AND CONCLUSIONS

OLAP mining integrates on-line analytical processing with data mining which substantially enhances the power and flexibility of data mining and makes mining an interesting exploratory process.

In this paper, we discussed the principles and some implementation techniques of OLAP mining by taking the OLAP mining functions of the DBMiner system as running examples.

With multiple data mining functions available, one question which naturally arises is how to determine which data mining function is the most appropriate one for a specific application. To select an appropriate data mining function, one needs to be familiar with the application problem, data characteristics, and the roles of different data mining functions. Sometimes one needs to perform interactive exploratory analysis to observe which function discloses the most interesting features in the database. Therefore, the building of exploratory analysis tools and the construction of an application-oriented semantic layer are two important solutions. OLAP mining provides an exploratory analysis tool, however, further study should be performed on the *automatic selection* of data mining functions for particular applications.

Another popularly posed question is how these data mining techniques are different from the set of existing statistical data analysis tools. Data mining is the confluence of multiple disciplines, including database systems, data

warehouses, statistics, visualization, machine learning, and information science. Previous work on statistics has provided some foundational work for data mining. Most effective data mining systems are extending the power of database systems or data warehouse systems and re-examining the methods studied in statistics, visualization, and machine learning. Many methods developed in data mining systems are novel and scalable, and integrates well with existing database systems, which are quite different from existing statistical analysis tools and machine learning packages. Thus, data mining forms a new direction in the research into the methods for the analysis of data and knowledge in large databases and data warehouses.

We are currently working on the further enhancement of the power and efficiency of OLAP mining of DBMiner for exploratory data mining, including the improvement of system performance and rule discovery quality for the existing functional modules, and the development of techniques for mining new kinds of rules, especially on time-related data, and visual data mining.

Another important task for future study is the extension of OLAP mining techniques towards advanced and/or special purpose database systems, including extended-relational, object-oriented, text, spatial, temporal, multimedia, and heterogeneous databases and Internet information systems. We will report the progress on OLAP mining of complex types of structured, semi-structured and nonstructured data in the future.

REFERENCES

Agarwal, S., Agrawal, R., Deshpande, P. M., Gupta, A., Naughton, J. F., Ramakrishnan, R. & Sarawagi, S. (1996), On the computation of multidimensional aggregates, *in* 'Proc. 1996 Int. Conf. Very Large Data Bases', Bombay, India, pp. 506–521.

Agrawal, R. & Srikant, R. (1994), Fast algorithms for mining association rules, *in* 'Proc. 1994 Int. Conf. Very Large Data Bases', Santiago, Chile, pp. 487–499.

Chaudhuri, S. & Dayal, U. (1997), 'An overview of data warehousing and OLAP technology', *ACM SIGMOD Record* **26**, 65–74.

Ester, M., Kriegel, H.-P., Sander, J. & Xu, X. (1996), A density-based algorithm for discovering clusters in large spatial databases, *in* 'Proc. 2nd Int. Conf. on Knowledge Discovery and Data Mining (KDD'96)', Portland, Oregon, pp. 226–231.

Fayyad, U. M., Piatetsky-Shapiro, G., Smyth, P. & Uthurusamy, R. (1996), *Advances in Knowledge Discovery and Data Mining*, AAAI/MIT Press.

Gray, J., Chaudhuri, S., Bosworth, A., Layman, A., Reichart, D., Venkatrao, M., Pellow, F. & Pirahesh, H. (1997), 'Data cube: A relational aggregation operator generalizing group-by, cross-tab and sub-totals', *Data Mining and Knowledge Discovery* **1**, 29–54.

Han, J. & Fu, Y. (1995), Discovery of multiple-level association rules from large databases, *in* 'Proc. 1995 Int. Conf. Very Large Data Bases', Zurich, Switzerland, pp. 420–431.

Han, J. & Fu, Y. (1996), Exploration of the power of attribute-oriented induction in data mining, *in* U. Fayyad, G. Piatetsky-Shapiro, P. Smyth & R. Uthurusamy, eds, 'Advances in Knowledge Discovery and Data Mining', AAAI/MIT Press, pp. 399–421.

Han, J., Chiang, J., Chee, S., Chen, J., Chen, Q., Cheng, S., Gong, W., Kamber, M., Liu, G., Koperski, K., Lu, Y., Stefanovic, N., Winstone, L., Xia, B., Zaiane, O. R., Zhang, S. & Zhu, H. (1997), DBMiner: A system for data mining in relational databases and data warehouses, *in* 'Proc. CASCON'97: Meeting of Minds', Toronto, Canada, pp. 249–260.

Jain, A. K. & Dubes, R. C. (1988), *Algorithms for Clustering Data*, Printice Hall.

Kamber, M., Han, J. & Chiang, J. Y. (1997), Metarule-guided mining of multi-dimensional association rules using data cubes, *in* 'Proc. 3rd Int. Conf. on Knowledge Discovery and Data Mining (KDD'97)', Newport Beach, California, pp. 207–210.

Mehta, M., Agrawal, R. & Rissanen, J. (1996), SLIQ: A fast scalable classifier for data mining, *in* 'Proc. 1996 Int. Conference on Extending Database Technology (EDBT'96)', Avignon, France.

Ng, R. & Han, J. (1994), Efficient and effective clustering method for spatial data mining, *in* 'Proc. 1994 Int. Conf. Very Large Data Bases', Santiago, Chile, pp. 144–155.

Quinlan, J. R. (1993), *C4.5: Programs for Machine Learning*, Morgan Kaufmann.

Sarawagi, S. (1997), 'Indexing OLAP data', *Bulletin of the Technical Committee on Data Engineering* **20**, 36–43.

Srikant, R. & Agrawal, R. (1995), Mining generalized association rules, *in* 'Proc. 1995 Int. Conf. Very Large Data Bases', Zurich, Switzerland, pp. 407–419.

Zhang, T., Ramakrishnan, R. & Livny, M. (1996), BIRCH: an efficient data clustering method for very large databases, *in* 'Proc. 1996 ACM-SIGMOD Int. Conf. Management of Data', Montreal, Canada, pp. 103–114.

Zhao, Y., Deshpande, P. M. & Naughton, J. F. (1997), An array-based algorithm for simultaneous multidimensional aggregates, *in* 'Proc. 1997 ACM-SIGMOD Int. Conf. Management of Data', Tucson, Arizona, pp. 159–170.