

ABSTRACT: Scientific databases contain large amounts of interrelated information. This information is often stored in relational databases with hundreds of tables and thousands of rows per table. Clustering is an effective way to reduce the information-overhead associated with finding information among these tables, allowing the user to browse through the clusters as well as the individual tables. In this paper, we compare the use of two adaptive algorithms (genetic algorithms, and neural networks) in clustering the tables of a scientific database. These clusters allow the user to index into this overwhelming number of tables and find the needed information quickly. We cluster the tables based on the user's queries and not on the content of the tables, thus the clustering reflects the unique relationships each user sees among the tables. The original database remains untouched, however each user will now have a personalized index into this database.

1 Introduction

Relational databases are designed to deal with limited ranges of data on specific topics. The form of the data is known ahead of time and the database tables and their relationships are clearly defined before the data is entered. Scientific databases [9] contain a much larger amount of data on many different, but related topics. The form of the data is not known ahead of time as the data is collected by investigators from a wide range of disciplines. Scientific data is often stored using relational databases. This huge amount of interrelated data is then forced into the rigid table structure of a relational database which can not adequately model the necessary relationships [7].

Scientific databases are accessed by users from a wide range of disciplines, mostly unfamiliar with databases and their associated query languages. These users need to search for specific pieces of data quickly, and browse through related information to see if it is of value to them. They need to relate information from different tables in the database.

When an interface to a scientific database is designed, its creator can either impose no structure on the database - there is just a large pool of tables; or impose a generic structure on the database - a hierarchical menu system allowing the user to move through an ordering of the tables.

Much of the difficulty in accessing data in scientific databases comes from the enormous amount of data that is involved; but the organization of this data is also a major problem [2]. Users from various scientific communities see

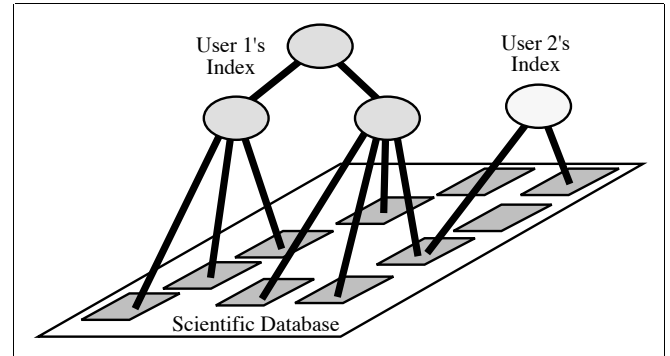


Figure 1: Hierarchical Index into the Scientific Database

different relationships between sets of data. Certain information is important to certain investigators and certain information is not. Typically a generic interface is provided to the data. This gives researchers from all backgrounds a way to access the data, but each of the researchers must conform to this generic structuring of the data.

Instead of forcing the users to conform to the structure of the database, we can mold the structure of the database to the needs of the individual users and thereby reduce their confusion when interacting with the database. Our solution is to cluster the tables of a scientific database based on the user's queries. The user of the database needs the benefits of clustering so we should involve the user in the clustering process. The content of the tables is irrelevant, only the fact that the user sees a relationship between the tables is important. This flexibility will be increasingly important as the size, breadth, and accessibility of scientific databases increases.

Ioannidis, et al [5] developed a graphical interface for the management of scientific experiments and data using the Object-Oriented data model MOOSE. The system makes large schemas more manageable by allowing the user to hide parts of the schema, collapse sections of the schema into nodes, use reference nodes to eliminate long arcs. This approach has several shortcomings: The users may not know what data is available, and may not know enough about the domain to make appropriate choices. It give users a variety of choices without sufficient descriptive material to make that choice.

We collect data on the user's queries. This information is then given to a genetic algorithm or to a neural network which partitions the tables of the scientific database into a hierarchy of clusters. The genetic algorithm and neural

network generate very good clusterings much faster than a deterministic exhaustive search algorithm. Each user now has the choice of browsing through the existing generic interface, or using a hierarchy of clusters as an index to move quickly to the appropriate tables in the database. The original scientific database remains untouched, however each user now has a personal index into this document. See Figure 1.

As new users begin to work with this database, they can choose from the existing clusterings. A new user can choose to look at the database from a biologist’s point of view, or a climatologist’s point of view. This gives each user a starting point nearer their own needs than the generic interface. Experienced users will also be able to access their data faster and more conveniently because the database will adapt to them.

Section 2 discusses how we collect information from the user. Section 3 discusses clustering in scientific databases. Sections 4 and 5 respectively discuss how we use a genetic algorithm and neural network to create the hierarchy of clusters. Section 6 discusses our implementation of both clustering methods. Section 7 discusses uses for this clustering. Finally, Section 8 gives our conclusions and plans for future work.

2 Information Collection

Users retrieve data from a relational database using a query language such as SQL. We monitor the user’s queries, and store information about the tables accessed in a list.

Each line of the list contains information on a single table in the database. This line lists all of the columns that have been displayed for that table, and all the columns that have been used to connect this table to other tables in the database. The format of each line of the list is shown below:

$$\text{table}_i : \{\text{selectCol}_i, \}^* : \{\text{joinCol}_i\text{-joinCol}_j @ \text{table}_j, \}^*.$$

where:

- table_i is the table we are interested in.
- selectCol_i s are the columns in table_i that have been selected.
- joinCol_i s are the columns in table_i that have been joined with columns joinCol_j of table table_j

Three small sample SQL queries are converted into their list form as shown in Figure 2.

When a user begins to search for information in the database they either choose to work with an existing list or create a new empty list. The user can give this new empty list an appropriate name. As the user works with the database, the list keeps track of the tables and columns the user is accessing through their queries.

The mechanism for creating the query (straight query commands, a graphical query language, etc.) is unimportant. Eventually the query in the form of tables and columns is given to the database and this query updates the listing.

Query 1	List 1
SELECT $A.\phi, B.\chi, C.\epsilon$ FROM A, B, C WHERE $A.\alpha = B.\beta$ AND $A.\alpha = C.\gamma$;	$A:\phi, \quad : \alpha\text{-}\beta @ B, \alpha\text{-}\gamma @ C, .$ $B:\chi, \quad : \beta\text{-}\alpha @ A, .$ $C:\epsilon, \quad : \gamma\text{-}\alpha @ A, .$
Query 2	List 2
SELECT $C.\rho, C.\omega, D.\lambda$ FROM C, D WHERE $C.\rho = D.\sigma$;	$C:\rho, \omega, \quad : \rho\text{-}\sigma @ D, .$ $D:\lambda, \quad : \sigma\text{-}\rho @ C, .$
Query 3	List 3
SELECT $A.\alpha$ FROM A, B, C WHERE $A.\delta = B.\beta$ AND $B.\beta = C.\rho$;	$A:\alpha, \quad : \delta\text{-}\beta @ B, .$ $B: \quad : \beta\text{-}\delta @ A, \beta\text{-}\rho @ C, .$ $C: \quad : \rho\text{-}\beta @ B, .$

Figure 2: Converting Queries into Lists

A user can create one list for all of their queries, or create multiple lists where each concept the user is interested in is given its own list. The lists can be unioned together. A user can keep their lists private or share them with other users of the scientific database. A list can therefore be one of a user’s lists, the union of a user’s lists, the union of a group’s lists, the union of a site’s lists, or the union of all the lists available on the system.

We have used a similar technique to collect information on a user’s browsing through a hypertext [10] document to cluster the nodes based on the user’s usage patterns [6]. The scientific database user, like a hypertext user, is faced with an overwhelming amount of interrelated information presented in a generic way. Clustering allows us to add a personal interface layer on top of the generic interface reducing information overhead and speeding up access to the underlying data.

3 Clustering

As scientific databases grow larger it becomes more and more important to cluster their tables. Clustering reduces information overhead. It allows for higher level concepts (groups of tables), allows for the breaking of a single large scientific database into appropriate modules, allows for views over the database, and allows the user to make changes at the cluster level without affecting the scientific database itself. Clustering has been found to be a very effective means of reducing information overhead in other large information systems with related components such as hypertext systems [1, 4].

By monitoring the user’s progress through the database we can cluster the scientific database based on the connections the user makes between the tables. Those tables that the user has found important are made more visible, and more easily accessible. In commonly used databases this will give the user quick access to important blocks of information, and several starting points for starting new searches.

Table	List 1 Percent	List 2 Percent	List 3 Percent	Neighbouring Tables
A	50	0	50	BC
B	50	0	50	AC
C	33	33	33	ABD
D	0	100	0	C

Figure 3: Converted Inputs to the Clustering Algorithms

The clusters can be shared between users. This gives new users a choice of several ways to view an unfamiliar database, and gives regular users a way to see the database from different perspectives.

We use a genetic algorithm, and a neural network to delineate the clusters. Finding the optimal clustering of the tables is equivalent to checking every partition of the set of tables. This is an exponential problem so it is too costly to find the optimal clustering for a database with more than a few tables. Using an adaptive algorithm such as genetic algorithms or neural networks we can find a very good clustering within a reasonable amount of time.

The lists that have been created by the information gathering process described in Section 2 are used to cluster the tables of the scientific database. The number of times each table is used in each list, divided by the total over all the lists is used to generate usage percentages. For each table, every joined table in each of the lists is added onto the roster of neighbouring tables.

The converted versions of the lists created in Figure 2 are shown in Figure 3. Since Table A was used once in List 1, and once in List 3, Lists 1 and 3 have usage percentages of 50% while List 2 has a usage percentage of 0%. Table A was joined with tables B and C in the queries so tables B, and C are neighbours of table A. The usage percentages and the neighbour roster are used to cluster the tables.

4 Genetic Algorithm Clustering

Genetic algorithms accept their input coded as a finite length string (or chromosome) over a finite alphabet. Each of the elements in the string is a gene, and each gene has an allele value (one of the values from the finite alphabet.) Several of these chromosomes form a population.

The population changes from one generation to the next through reproduction, crossover, and mutation. The number of chromosomes in each iteration remains constant, but their fitness should generally improve. The more fit a chromosome is, the more likely it is to reproduce and survive until the next generation. The fitness of a chromosome is determined by a payoff function which depends on the allele values of each chromosome.

Each of the chromosomes in the population is initially given random allele values. The more fit chromosomes of the previous generation reproduce to produce the chromosomes of the next generation. Sometimes when two chromosomes reproduce, they crossover. When chromosomes A

and B crossover to form chromosomes C and D, a random crossover percentage is generated. That percentage of chromosome A’s genes are copied into identical positions in C. The remaining genes in chromosome A are copied into D. Genes from chromosome B fill in the remaining positions in C and D. The result of the crossover may be a more fit chromosome, or a less fit chromosome.

Crossover only mixes existing allele values at each gene. Mutation causes random changes in the allele values, allowing for more (uncontrolled) variation. The result of the mutation may be a more fit chromosome, or a less fit chromosome. For a more thorough introduction to genetic algorithms see [3].

In our approach, each chromosome represents a set of clusters; each gene represents a table in the scientific database; each allele value represents a cluster. Genes with the same allele value are in the same cluster. Each table is a member of at most one cluster, and the number of clusters varies from 1 to n where n is the total number of tables being clustered. Each table can form its own unique cluster in the chromosome, if necessary.

The payoff function is composed of two separate partial payoff functions: weight similarity, and neighbourliness. Each partial payoff function’s values range from 0.0 to 10.0.

Weight similarity, computed by Equation 1, promotes clusters containing tables with similar weight percentages, indicating similar usage patterns. For every cluster, the difference between each table percentage and the average percentage of the cluster is computed and scaled into the range 0 to 10. The closer the percentages are, the smaller this ratio will be, and the larger the overall payoff will be. Weight similarity pushes the genetic algorithm to create smaller clusters.

$$10 \times (1 - \min(\frac{\sum_{tables} \sum_{lists} |table\% - aveCluster\%|}{\# of tables}, 1)) \quad (1)$$

Neighbourliness, computed by Equation 2, promotes clusters containing neighbouring tables. The roster of neighbours shows which tables have been joined. For each pair of tables that are connected by a join, at least one user must feel these two tables are related because that user joined them. Each connection in the input lists either connects two tables in the same cluster, or two tables in different clusters. Neighbourliness sums up the number of connections between tables in the same cluster and divides the total by the total number of tables in the lists. Neighbourliness pushes the genetic algorithm to create bigger clusters.

$$10 \times \frac{\sum_{clusters} \# of joins between tables in cluster}{\# of joins} \quad (2)$$

Neighbourliness tries to bring all the connected tables together. Weight similarity tries to isolate tables with similar usage patterns. Together, this pulling together and pushing apart generates the clusters. Equation 3 shows how the two partial payoff functions are combined.

$$payoff = neighbourliness \times weightSimilarity^2 \quad (3)$$

This results in an overall payoff function value that ranges from 0 to 1000.

Each table is initially put into its own cluster. This generates better clusters than randomly generating initial clusters. The genetic algorithm is started. After each new generation is created the most fit chromosome of the previous generation replaces the least fit chromosome of the new generation. This preserves the best set of clusters from one generation to the next. The genetic algorithm stops when the value of the most fit chromosome remains constant, or the number of generations reaches a set maximum value.

If there is more than one resulting cluster, these clusters are given back to the genetic algorithm. The weight percentages, and neighbours of the components of each cluster are used to generate the weight percentage, and neighbours for the cluster. Each cluster is now treated like a table. The clusters are clustered forming a hierarchy until one global cluster remains.

5 Neural Network Clustering

A neural network is a distributed network of small processing elements. This large number of simple processing units is connected together forming a miniature version of the human brain. Neural networks have had a good deal of success in pattern recognition and learning by example.

We are using a self organizing feature map, a neural network developed by Teuvo Kohonen [8]. A self organizing feature map is a two dimensional array of processing elements. Each element contains a weight vector. The input to the self organizing feature map is a set of weight vectors. The self organizing feature map forms itself into a topological ordering of the input data through unsupervised learning.

The weight vector W_{ij} of each element of the N by N array is initialized with random values. An input vector I is chosen at random. The neuron whose weight vector W_{ij} is closest to I is found ($\min |W_{ij} - I|$). This neuron and all the neurons within a certain neighbourhood (N_c by N_c) of this neuron have their weights adjusted, bringing the weights closer to I. Random inputs are repeatedly provided while the amount of adjustment (α), and the neighbourhood shrink in a form of simulated annealing. The amount of adjustment and the size of the neighbourhood eventually become negligible and the learning ceases.

When the learning stops, the network has organized the inputs into a two dimensional array. For each input there is one processing element with the closest matching weight vector. Labeling each neuron with the table that it matches most closely gives the topological ordering of the inputs. For a more thorough introduction to neural networks see [11].

In our clustering algorithm the weight vectors contain the usage percentages for the various usage lists being used to cluster the database. If there are k lists involved in the

Abbreviated Table Name	Table Description
AEROLOG	measurements from surface flux group
AIR_FLUX	aircraft flux from flights over the konza ave from ncar's pams, army corp's dcps
AMS_87	ave from ncar's pams, army corp's dcps
AMS_89	ave from ncar's pams, army corp's dcps
AMS_STAT	# of reports from each ams station
BIOMASS	plant biomass weight, nitrogen content
BRUT	actual radiosonde data observations
CLOUD	cloud estimates from liverpool cameras
FIFE_SITE	reference info. on the collection sites
GRAV	soil moisture readings at 25,75,150mm
NEUTRON	soil moisture with 200cm neutron probe
RAD_FLUX	measurements from surface flux group
RAIN_DAY	daily rainfall data by site and date
SOIL_PROP	soil properties measured historically
SOIL_GAS	no2 flux, co2 from soil respiration
VEG_SPEC	species composition data by site, date
WIND_PRO	noaa lidar wind profile data

Figure 4: FIFE Table Descriptions

clustering, the weight vector will contain k values. The size of the array $N = \lceil \sqrt{\text{number of tables}} \rceil$. Each table can form its own unique cluster in the array, if necessary.

The amount of adjustment, and the neighbourhood are adjusted as follows:

$$\alpha_0 = 1$$

$$\alpha_{t+1} = \begin{cases} \alpha_t - 0.001 & \text{if } \alpha_t \geq 0.7 \\ \alpha_t - 0.0001 & \text{otherwise} \end{cases}$$

$$N_{c_t} = N \times \alpha_t$$

The weigh vectors are adjusted as follows:

$$W_{ij_{t+1}} = \begin{cases} W_{ij_t} + \alpha_t \times (I - W_{ij_t}) & \text{if node is within } N_{c_t} \\ W_{ij_t} & \text{otherwise} \end{cases}$$

For each iteration in the learning process, one table is chosen at random. The neuron with the most similar set of weights is identified. That neuron and all the neurons within the current neighbourhood have their weights adjusted bringing them closer to the weights of the chosen table. Each neighbour of the chosen table is taken in turn. The neuron with the most similar set of weights to the neighbouring table is identified, and its weights are adjusted to bring it closer to the original table chosen. The self organizing feature map orders itself with regards to similarity in the weight percentages and the neighbouring tables.

6 Implementation

We tested our approach on NASA's FIFE scientific database.

The objective of the ISLSCP (International Satellite Land Surface Climatology Project) is to develop techniques to determine surface climatology from satellite observations. FIFE (First ISLSCP Field Experiment) was undertaken at a 15km by 15km square site near Manhattan, Kansas in 1987

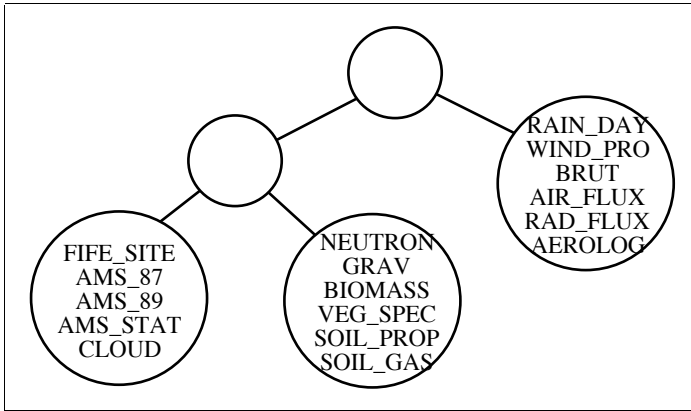


Figure 5: Clusters produced by Genetic Algorithm

and 1989. Its purpose was to gather enough data to allow the creation and testing of models to develop these techniques [12]. 120 gigabytes of data was collected (300 megabytes textual data, the rest images.) The textual data fills 100 tables in a relational database.

Five queries utilizing 17 tables were made to the database showing the kind of queries that researchers in different, but related fields would create. Each of these queries was stored in its own appropriately named list. The tables which made up those queries are shown in Figure 4. The tables used in each of the five queries are shown below:

1. Rainfall: FIFE_SITE, AMS_87, AMS_89, AMS_STAT
2. Meteorological: FIFE_SITE, AMS_87, AMS_89, AMS_STAT, RAIN_DAY, CLOUD
3. Atmospheric: FIFE_SITE, WIND_PRO, BRUT, AIR_FLUX, RAIN_DAY, RAD_FLUX, AEROLOG
4. Surface Biophysical: FIFE_SITE, BIOMASS, SOIL_PROP, SOIL_GAS, CLOUD, VEG_SPEC
5. Soil Moisture: FIFE_SITE, CLOUD, NEUTRON, GRAV, BIOMASS, SOIL_PROP

The genetic algorithm takes 25 seconds to return the clusters shown in Figure 5. The neural network takes 70 seconds to return the clusters shown in Figure 6.

The genetic algorithm and neural network create similar clusterings of the database tables. The genetic algorithm clusters the database into 3 clusters, which are clustered into 2 clusters which are then clustered into one global cluster. The neural network clusters the databases into 8 clusters which are clustered into one cluster. The three clusters at the top of the neural network form the central cluster in the genetic algorithm. The two clusters at the bottom right of the neural network form the rightmost cluster in the genetic algorithm. The three clusters in the center and lower left of the neural network form the leftmost cluster in the genetic algorithm.

The genetic algorithm generates a strict hierarchy with an arbitrary number of levels where the relationships between the clusters are shown by their relative positions in

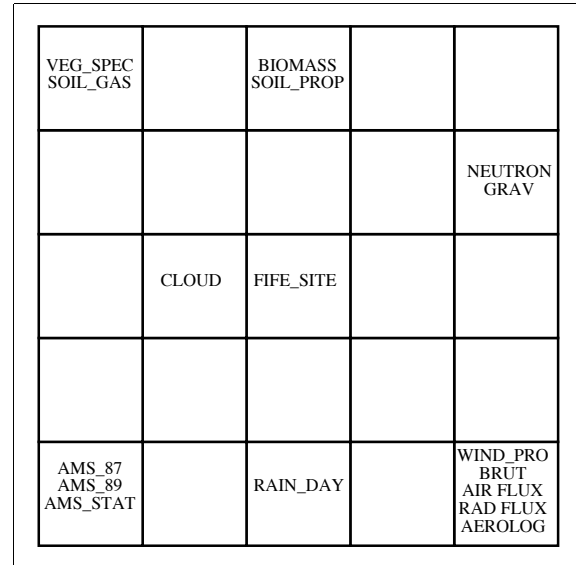


Figure 6: Clusters produced by Neural Network

the hierarchy. The neural network generates only two levels of clustering but the overview shows more subtle relationships between the clusters. The genetic algorithm generates deeper, more structured indices than the neural network. The neural network gives a better overall picture of how clusters are related.

The neural network clustering can help a new user see a global picture of the scientific database. With the neural network clustering, tables used most generally will be found near the center of the array. Tables with more specific uses will be found towards the corners. This can be seen in Figure 6 where the FIFE_SITE table used in all five queries was placed in the center of the array while the VEG_SPEC table used in only one query was placed in the upper left corner.

The genetic algorithm clustering can help a user familiar with the system categorize clusters into a hierarchical relationship. General concepts are found in the interior nodes at the top of the hierarchy while more specific concepts are found at the bottom of the hierarchy, near the actual tables being clustered.

The two algorithms have comparable execution times. With inputs of less than 14 tables our genetic algorithm, on average, found a clustering better than 99% of the possible clusterings. It was impractical to do comparisons with more than 13 input tables due to the running time of the deterministic program. For a comparison of running times of the deterministic, genetic algorithm, and neural network clustering algorithms see Figure 7. The time for running the neural network is directly related to the size of the array. As the array size is based on the $\sqrt{\text{number of tables}}$, the time complexity of the neural network is a step function.

Clustering will be an ongoing process in the database. New users will begin using the database, current users will return to search for new data. Perhaps even new tables will be added to the database. More and more relationships will connect the tables together, and more tables will be brought

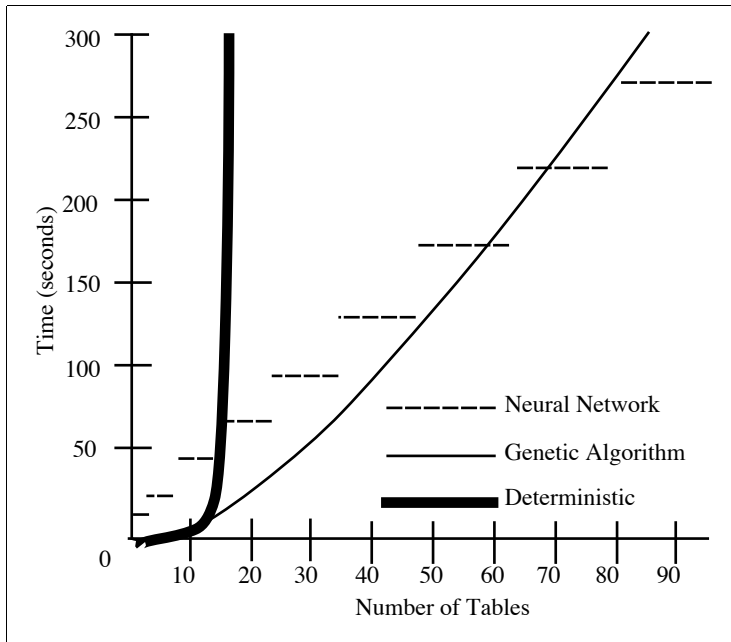


Figure 7: Comparison of Different Clustering Algorithms

into the clustering. These new relationships may be very different from the older relationships so the system should dynamically support both the new and current users.

Totally re-clustering the tables will allow the index to show the most effective clustering based on all of the users' previous experiences. This will be very useful to new users, but may be confusing to current users when the clustering patterns change, possibly dramatically. Instead the clustering can be performed incrementally. The current clusters can be used as an input to both the genetic algorithm and the neural network, instead of starting the clusters from scratch. This will allow the clusters to change slowly, giving the current users a common framework and adding in each new user's contributions.

More than one set of indices can be maintained in the system. New users to start with the most up-to-date indexing information, while current users can access the database using their familiar clusters. However the clusters are maintained, the underlying database remains untouched.

7 Uses

The clusters that have been generated can be used for several purposes. They can aid both the user and the database administrator.

When a user wishes to start a new query, they can look through the existing set of topics that have been queried. One of these topics may match up with their needs giving them an excellent starting point for their queries. They may find that their area of interest straddles two existing queries. Using the existing lists the user can reduce their search time in the database as shown in Figures 8 and 9.

Figure 8 shows how the clusters in Figure 6 can be high-

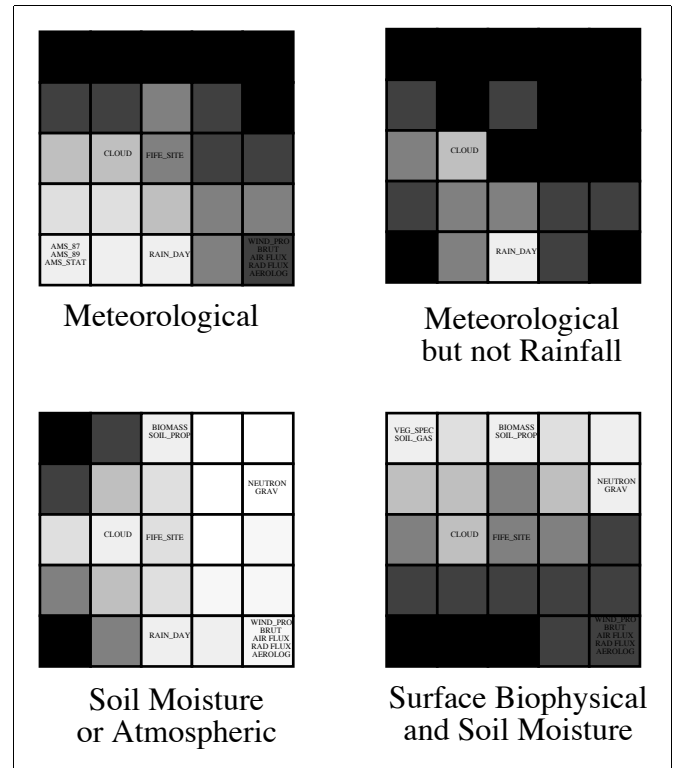


Figure 8: Sample Index Queries using Neural Network

lighted. Each of the arrays in Figure 8 is in the same orientation as the larger (and more readable) Figure 6. The user can choose to see the tables important to someone interested in Meteorological information, or someone interested in Meteorological data but not Rainfall data. The user might want information related to Soil Moisture or Atmospheric conditions, or they may want information that combines Surface Biophysical data and Soil Moisture data.

By selecting the appropriate queries the clusters highlight the appropriate areas of the database. Important areas are shown in white; unimportant areas are shown in black. The brightness of the region shows how relevant the tables in that region are to the suggested query areas. Since the neural network forms a topological map of the input lists, these input lists allow us to find the literal peaks and valleys of interest.

As well as highlighting the important tables, the clustering diagram also shows relationships between the highlighted tables, and how general or specific the highlighted tables are by their location in the array. Thus it gives much more information to the user than simply highlighting table names based on the usage percentages directly, or using a general overview diagram of the database.

Figure 9 shows how the clusters in Figure 5 can be highlighted. As with the neural network version, the user can choose combinations of the existing queries to see the tables that are of personal importance.

Using the clustering information provided by the users, the database administrator can find which tables are in common usage among which groups. For example, all the inves-




Figure 9: Sample Index Query using Genetic Algorithm

tigators would require access to the most commonly used tables but perhaps only the geologists are accessing certain tables, and the climatologists are the exclusive users of another set of tables. Using the clusters the administrator can determine which tables an investigator will need based on their areas of interest, because those are the tables that investigators with similar interests have used in the past.

8 Conclusions and Future Work

In this paper we have proposed, implemented, and compared two adaptive clustering algorithms for scientific databases. These algorithms hierarchically cluster tables according to each user's view of the scientific database, rather than the database administrator's view of the data. This hierarchy of clusters gives each user a personal index into the scientific database. New and experienced users can use the existing indices to speed up their access to the database.

Both the genetic algorithm and the neural network create clusters within a reasonable amount of time. The genetic algorithm generates a strict multi-level hierarchy of clusters. The neural network generates a two dimensional map of the clusters. Both algorithms allow the user to view the scientific database in a personalized way.

We are currently enhancing our implementation in the following ways: 1. Improving cluster visualization using 3d graphics, 2. Applying this approach to other database systems (e.g. object oriented scientific databases), 3. Using the stored list information to suggest appropriate join and display columns, 4. Decreasing clustering time through parallelism, and 5. Considering alternative clustering techniques.

Acknowledgements

We gratefully acknowledge the assistance of John Norman of the University of Wisconsin, and Don Strebel of NASA for their insight into the FIFE database.

References

- [1] Botafogo, R., Schneiderman, B. "Identifying aggregates in hypertext structures." In *Proceedings of the Hypertext '91 Conference*, (San Antonio, Texas, Dec 15-18, 1991), 63-74.
- [2] French, J., Jones, A., Pfaltz, J. "Summary of the final report of the nsf workshop on scientific database management." In *SIGMOD RECORD*, vol 19, no. 4, (Dec. 1990), 32-40.
- [3] Goldberg, D. "Genetic algorithms in search, optimization, and machine learning." Addison-Wesley, New York (1989).
- [4] Hara, Y., Keller, A., Wiederhold, G. "Implementing hypertext database relations through aggregations and exceptions." In *Proceedings of the Hypertext '91 Conference*, (San Antonio, Texas, Dec 15-18, 1991),
- [5] Ioannidis, Y., Livny, M., Haber, E. "Graphical user interfaces for the management of scientific experiments and data." In *SIGMOD RECORD*, vol 21, no 1, (March, 1992), 47-53.
- [6] Johnson, A., Fotouhi, F. "Automatic touring in a hypertext system." In *Proceedings of the IEEE International Phoenix Conference on Computers and Communication*, (Tempe, Arizona, Mar 23-26, 1993), 524-530.
- [7] Kim, W., "Object-Oriented Approach to Managing Statistical & Scientific Databases" In *Proceedings of the Fifth International Conference on Statistical & Scientific Database Management*, (Charlotte, N.C., April 1990).
- [8] Kohonen, T. "Self-organizing feature map, the." In *Proc. IEEE*, vol 78, no. 9, (September, 1990), 1464-1480.
- [9] Michalewicz, Z. (ed) "Statistical and scientific databases." Ellis Horwood, (1991).
- [10] Nielsen, J. "Hypertext and hypermedia." Academic Press, San Diego (1990).
- [11] Nielsen, R. "Neurocomputing." Addison-Wesley, New York (1990).
- [12] Strebel, D., Newcomer, J., Ormsby, J. "Data management in the FIFE information system." In *Proceedings of International Geoscience and Remote Sensing Symposium*, (Vancouver, B.C., Canada, July, 1989), 42-45.