

# WebACE: A Web Agent for Document Categorization and Exploration

Eui-Hong (Sam) Han, Daniel Boley, Maria Gini,  
Robert Gross, Kyle Hastings, George Karypis,  
Vipin Kumar, Bamshad Mobasher, and Jerome Moore  
Department of Computer Science and Engineering,  
University of Minnesota

## Abstract

We propose an agent for exploring and categorizing documents on the World Wide Web based on a user profile. The heart of the agent is an automatic categorization of a set of documents, combined with a process for generating new queries used to search for new related documents and filtering the resulting documents to extract the set of documents most closely related to the starting set. The document categories are not given a-priori. The resulting document set could also be used to update the initial set of documents. We present the overall architecture and describe two novel algorithms which provide significant improvement over traditional clustering algorithms and form the basis for the query generation and search component of the agent.

## 1 Introduction

The World Wide Web is a vast resource of information and services that continues to grow rapidly. Powerful search engines have been developed to aid in locating unfamiliar documents by category, contents, or subject. Relying on large indexes to documents located on the Web, search engines determine the URLs of those documents satisfying a user's query. Often queries return inconsistent search results, with document referrals that meet the search criteria but are of no interest to the user.

While it may not be currently feasible to extract in full the meaning of an HTML document, intelligent software agents have been developed which extract semantic features from the words or structure of an HTML document. These extracted features are then employed to classify and categorize the documents. Clustering offers the advantage that *a priori* knowledge of categories is not needed, so the categorization process is unsupervised. The results of clustering could then be used to automatically formulate queries and search for other similar documents on the Web, or to organize bookmark files, or to construct a user profile.

In this paper, we present WebACE, an agent for document categorization and exploration that operates on Web documents. A novel part of the paper is the description of two new clustering algorithms based on graph partitioning, that provide a significant improvement

in performance over traditional clustering algorithms used in information retrieval. Many traditional algorithms break down as the size of the document space, and thus the dimensionality of the corresponding feature space, increases. High dimensionality is characteristic of the type of information retrieval applications which are used to filter and categorize hypertext documents on the World Wide Web. In contrast, our partitioning-based algorithms do not rely on a specific choice of a distance function and do scale up effectively in a high dimensional space.

After a short description of the architecture of WebACE in Section 3, we describe the clustering algorithms in Section 4. The results obtained on a number of experiments using different methods to select sets of features from the documents show that partitioning clustering methods perform better than traditional distance based clustering. In Section 5 we show how to use words obtained from clusters of documents to generate queries for related documents on the Web.

## 2 Related Work

The heterogeneity and the lack of structure that permeates much of the information sources on the World Wide Web makes automated discovery, organization, and management of Web-based information difficult. Traditional search and indexing tools of the Internet and the World Wide Web such as Lycos, Alta Vista, WebCrawler, MetaCrawler, and others provide some comfort to users, but they do not generally provide structural information nor categorize, filter, or interpret documents.

In recent years these factors have prompted researchers to develop more intelligent tools for information retrieval, such as intelligent Web agents. The agent-based approach to Web mining involves the development of sophisticated AI systems that can act autonomously or semi-autonomously on behalf of a particular user, to discover and organize Web-based information. Generally, the agent-based Web mining systems can be placed into the following categories:

**Intelligent Search Agents** Several intelligent Web agents have been developed that search for relevant information using characteristics of a particular domain (and possibly a user profile) to organize and interpret the discovered information. For example, agents such as FAQ-Finder [HBML95], Information Manifold [KLSS95], and OCCAM [KW96] rely either on pre-specified and domain specific information about particular types of documents, or on hard coded models of the information sources to retrieve and interpret documents. Other agents, such as ShopBot [DEW96] and ILA [PE95], attempt to interact with and learn the structure of unfamiliar information sources. ShopBot retrieves product information from a variety of vendor sites using only general information about the product domain. ILA, on the other hand, learns models of various information sources and translates these into its own internal concept hierarchy.

**Information Filtering/Categorization** A number of Web agents use various information retrieval techniques [FBY92] and characteristics of open hypertext Web documents to automatically retrieve, filter, and categorize. For example, HyPursuit [WVS<sup>+</sup>96] uses semantic information embedded in link structures as well as document content

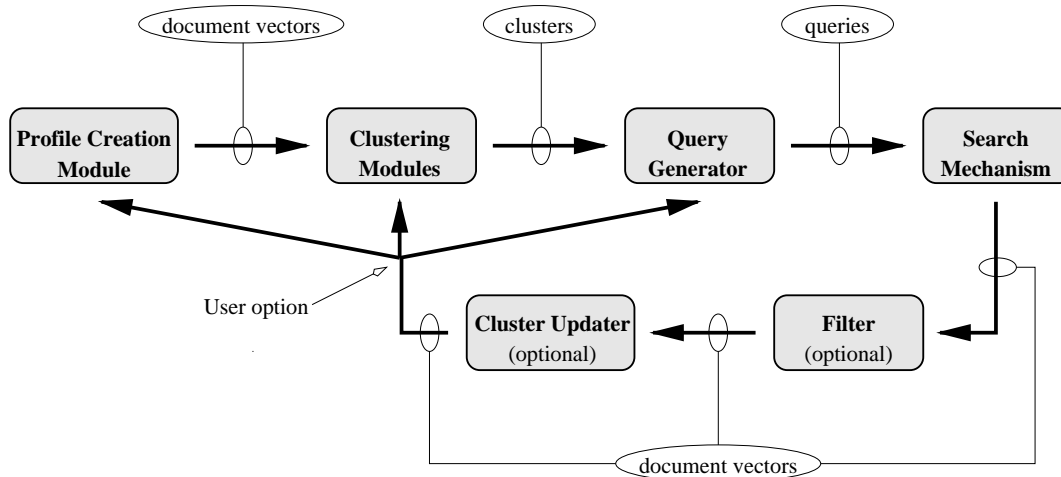


Figure 1: WebACE Architecture

to create cluster hierarchies of hypertext documents, and structure an information space. BO (Bookmark Organizer) [MS96] combines hierarchical clustering techniques and user interaction to organize a collection of Web documents based on conceptual information. Pattern recognition methods and word clustering using the Hartigan’s K-means partitional clustering algorithm are used in [WP97] to discover salient HTML document features (words) that can be used in finding similar HTML documents on the Web.

**Personalized Web Agents** Another category of Web agents includes those that obtain or learn user preferences and discover Web information sources that correspond to these preferences, and possibly those of other individuals with similar interests (using collaborative filtering). A few recent examples of such agents include WebWatcher [AFJM95], Syskill & Webert, and others. For example, Syskill & Webert [Ack97] utilizes a user profile and learns to rate Web pages of interest using a Bayesian classifier. Balabanovic [BSY95] uses a single well-defined profile to find similar web documents. Candidate web pages are located using best-first search. The system needs to keep a large dictionary and is limited to a single user.

### 3 WebACE Architecture

WebACE’s architecture is shown in Figure 1. As the user browses the Web, the profile creation module builds a custom profile by recording documents of interest to the user. The number of times a user visits a document and the total amount of time a user spends viewing a document are just a few methods for determining user interest [Ack97, AFJM95, BSY95]. Once WebACE has recorded a sufficient number of interesting documents, each document is reduced to a document vector and the document vectors are passed to the the clustering modules. WebACE uses two novel algorithms for clustering which can provide significant improvement in both run-time performance and cluster quality over traditional algorithms. These are described in Section 4.

After WebACE has found document clusters, it can use the clusters to generate queries and search for similar documents. WebACE submits the queries to the search mechanism and gathers the documents returned by the searches, which are in turn reduced to document vectors. These new documents can be used in a variety of ways. One option is for WebACE to cluster the new documents, filtering out the less relevant ones. Another is to update the existing clusters by having WebACE insert the new documents into the clusters. Yet another is to completely re-cluster both the new and old documents. Finally, the user can decide to add any or all of the new documents to his profile. The query generation methods and the algorithms for incrementally updating existing clusters are discussed in Section 5.

WebACE is implemented as a browser independent Java application. Monitoring the user's browsing behavior is accomplished via a proxy server. The proxy server allows WebACE to inspect the browser's HTTP requests and the resulting responses. Upon execution, WebACE spawns a browser and starts a thread to listen for HTTP requests from the browser. As the browser makes requests, WebACE creates request threads to handle them. This allows multi-threaded browsers the capability of having multiple requests pending at one time. The lifespan of these request threads is short, i.e. the duration of one HTTP request, Conversely, the browser listener thread persists for the duration of the application.

## 4 Clustering Methods

Existing approaches to document clustering are generally based on either probabilistic methods, or distance and similarity measures (see [FBY92]). Distance-based methods such as  $k$ -means analysis, hierarchical clustering [JD88] and nearest-neighbor clustering [LF78] use a selected set of words (features) appearing in different documents as the dimensions. Each such feature vector, representing a document, can be viewed as a point in this multi-dimensional space.

There are a number of problems with clustering in a multi-dimensional space using traditional distance- or probability-based methods. First, it is not trivial to define a distance measure in this space. Some words are more frequent in a document than other words. Simple frequency of the occurrence of words is not adequate, as some documents are larger than others. Furthermore, some words may occur frequently across documents. Techniques such as TFIDF [SM83] have been proposed precisely to deal with some of these problems.

Secondly, the number of all the words in all the documents can be very large. Distance-based schemes generally require the calculation of the mean of document clusters. If the dimensionality is high, then the calculated mean values do not differ significantly from one cluster to the next. Hence the clustering based on these mean values does not always produce very good clusters. Similarly, probabilistic methods such as Bayesian classification used in AutoClass [CS96], do not perform well when the size of the feature space is much larger than the size of the sample set. This type of data distribution seems to be characteristic of document categorization applications on the Web, such as categorizing a bookmark file. Furthermore, the underlying probability models usually assume independence of attributes (features). In many domains, this assumption may be too restrictive. It is possible to reduce the dimensionality by selecting only frequent words from each document, or to use some other method to extract the salient features of each document. However, the number of

features collected using these methods still tends to be very large, and due to the loss of some of the relevant features, the quality of clusters tends not to be as good.

Our proposed clustering algorithms which are described in this section are designed to efficiently handle very high dimensional spaces, and furthermore, they do not require the definition of ad hoc distance or similarity metrics. In contrast to traditional clustering methods, our proposed methods are linearly scalable, an advantage which makes these methods particularly suitable for use in Web retrieval and categorization agents. For our evaluation, we compare these algorithms to two well-known methods: Bayesian classification as used by AutoClass [CS96] and *hierarchical agglomeration clustering (HAC)* based on the use of a distance function [DH73].

AutoClass is based on the probabilistic mixture modeling [TSM85], and given a data set it finds maximum parameter values for a specific probability distribution functions of the clusters. The clustering results provide the full description of each cluster in terms of probability distribution of each attributes. The HAC method starts with trivial clusters, each containing one document and iteratively combines smaller clusters that are sufficiently “close” based on a distance metric. In HAC, the features in each document vector is usually weighted using the TFIDF function, which is an increasing function of the feature’s text frequency and its inverse document frequency in the document space.

## 4.1 Association Rule Hypergraph Partitioning Algorithm

In [HKKM97], a new method was proposed for clustering related items in transaction-based databases, such as supermarket bar code data, using association rules and hypergraph partitioning. This method first finds set of items that occur frequently together in transactions using association rule discovery methods [AMS<sup>+</sup>96]. These frequent item sets are then used to group items into hypergraph edges, and a hypergraph partitioning algorithm [KAKS97] is used to find the item clusters. The similarity among items is captured implicitly by the frequent item sets.

In document clustering, each document corresponds to an item and each possible feature corresponds to a transaction. A frequent item sets found using the association rule discovery algorithm corresponds to a set of documents that have a sufficiently large number of features in common. These frequent item sets are mapped into hyperedges in a hypergraph. A hypergraph [Ber76]  $H = (V, E)$  consists of a set of vertices  $V$  and a set of hyperedges  $E$ . A hypergraph is an extension of a graph in the sense that each hyperedge can connect more than two vertices. In this model, the set of vertices  $V$  corresponds to the documents, and each hyperedge  $e \in E$  corresponds to a set of related documents found. For example, if  $\{d_1, d_2, d_3\}$  is a frequent item set, then the hypergraph contains a hyperedge that connects  $d_1, d_2$  and  $d_3$ . The weight of a hyperedge is calculated as the average confidence [AMS<sup>+</sup>96] of all the association rules involving the related documents of the hyperedge. The confidence of an association rule involving documents like  $\{d_1, d_2\} \implies \{d_3\}$  is the conditional probability that a feature occurs in document  $d_3$  whenever it occurs in  $d_1$  and  $d_2$ .

Next, a hypergraph partitioning algorithm is used to partition the hypergraph such that the weight of the hyperedges that are cut by the partitioning is minimized. Note that by minimizing the hyperedge-cut we essentially minimize the relations that are violated by partitioning the documents into different clusters. Similarly, this method can be applied

to word clustering. In this setting, each word corresponds to an item and each document corresponds to a transaction.

This method uses the *A priori* algorithm [AMS<sup>+</sup>96] which has been shown to be very efficient in finding frequent item sets and HMETIS [KAKS97] which can partition very large hypergraphs (of size > 100K nodes) in minutes on personal computers.

An additional advantage of ARHP is that it can be used to filter out non-relevant documents while clustering a document space, and thus improving the quality of the document clusters. This filtering capability is mainly due to support criteria in the association rule discovery components of the algorithm. Depending on the support threshold, documents that do not meet support (i.e., documents that do not share large enough subsets of words with other documents) will be pruned. This feature is particularly useful for clustering large document sets which are returned by standard search engines using keyword queries.

## 4.2 Partitioning Using Principal Component Analysis

The Principal Component Method is based on the computation of the leading principal direction for a collection of documents and then cutting the collection of documents along a hyperplane resulting in two separate clusters. The algorithm is then repeated on each separate cluster. The principal direction, which can be thought of as the direction of maximum variance, is computed using the Singular Value Decomposition (SVD). Each document is represented by a column of word counts and all the columns are collected into a matrix, just as in Latent Semantic Indexing (LSI) [BDO95]. However, from this point the process differs from that in [BDO95] in that (a) we first scale the columns to have unit length to make the results independent of the document length, (b) we translate the collection of document columns so that their mean lie at the origin, (c) we compute only the single leading singular value with its associated left and right singular vectors. In LSI as described in [BDO95], the SVD is applied to the original untranslated matrix of word counts, and the first  $k$  singular values and associated vectors are retrieved, for some choice of  $k$ . This removes much of the noise present in the data, and also yields a representation of the documents of reduced dimensionality, reducing the cost and raising the precision of user queries.

In contrast, in our Principal Component algorithm we form the hyperplane through the mean (centroid) of the cluster of documents whose normal is in the principal direction, and use the hyperplane to separate the documents into two subclusters. Since we need only one singular value, a simple fast Lanczos-based algorithm can be used to compute it [GV96]. The scaling is limited to normalizing the columns to make the results independent of the document lengths. This scaling maintains the sparsity structure in the matrix representing the documents. In our examples, the fill in the document matrix (percentage of elements that were nonzero) was 3 to 4%. TFIDF scaling fills up the matrix completely making it impractical to carry out the SVD computations while not yielding much improvement in the clustering results, at least in our examples.

The Principal Component computes a root hyperplane, and then a child hyperplane for each cluster formed from the root hyperplane, and so on. The result is a binary tree of clusters defined by associated hyperplanes. The algorithm proceeds by splitting a leaf node into two children nodes using the leaf's associated hyperplane. The leaf to be split is selected based on a modified scatter value [DH73] consisting of the sum of squares of the distances

from each document in the cluster to the cluster mean.

### 4.3 Experimental Evaluation of Clustering

To compare our clustering methods with the more traditional algorithms, we selected 185 web pages in 10 broad categories: business capital (BC), intellectual property (IP), electronic commerce (EC), information systems (IS), affirmative action (AA), employee rights (ER), personnel management (PM), industrial partnership (IPT), manufacturing systems integration (MSI), and materials processing (MP). The pages in each category were obtained by doing a keyword search using a standard search engine. These pages were then downloaded, labeled, and archived. The labeling facilitates an entropy calculation and subsequent references to any page were directed to the archive. This ensures a stable data sample since some pages are fairly dynamic in content.

Word Set	Selection Criteria	Dataset Size	Comments
E1	All words	185x10536	We select all non-stop words (stemmed).
E2	All words with text frequency > 1	188x5106	We prune the words selected for E11 to exclude those occurring only once.
E3	Top 20+ words	185x1763	We select the 20 most frequently occurring words and include all words from the partition that contributes the 20th word.
E4	Top 20+ with text frequency > 1	185x1328	We prune the words selected for E3 to exclude those occurring only once.
E5	Top 15+ with text frequency > 1	185x1105	
E6	Top 10+ with text frequency > 1	185x805	
E7	Top 5+ with text frequency > 1	185x474	
E8	Top 5+ plus emphasized words	185x2951	We select the top 5+ words augmented by any word that was emphasized in the html document, i.e., words appearing in <TITLE>, <H1>, <H2>, <H3>, <I>, <BIG>, <STRONG>, or <EMPHASIZE> tags.
E9	Quantile filtering	185x946	Quantile filtering selects the most frequently occurring words until the accumulated frequencies exceed a threshold of 0.25, including all words from the partition that contributes the word that exceeds the threshold.
E10	Frequent item sets	185x499	We select words from the document word lists that appear in a-priori word clusters. That is, we use an object measure to identify important groups of words.

Table 1: Setup of experiments.

The word lists from all documents were filtered with a stop-list and “stemmed” using




Figure 2: Class distribution of AutoClass clusters.

Porter’s suffix-stripping algorithm [Por80] as implemented by [Fra92]. We derived 10 experiments (according to the method used for feature selection) and clustered the documents using the four algorithms described earlier. The objective of feature selection was to reduce the dimensionality of the clustering problem while retain the important features of the documents. Table 1 shows the feature selection methods that characterize various experiments.

Validating clustering algorithms and comparing performance of different algorithms is complex because it is difficult to find an objective measure of quality of clusters. We decided to use entropy as a measure of goodness of the clusters (with the caveat that the best entropy is obtained when each cluster contains exactly one document). For each cluster of documents, the class distribution of documents is calculated first. Then using this class distribution, the entropy of each cluster is calculated. When a cluster contains documents from one class only, the entropy value is 0.0 for the cluster and when a cluster contains documents from many different classes, then entropy of the cluster is higher. The total entropy is calculated as the weighted sum of entropies of the clusters. We compare the results of the various experiments by comparing their entropy across algorithms and across feature selection methods (Fig. 5). Figure 2 shows the class distribution of documents in each cluster of the best AutoClass result with the entropy value 2.05. Comparing this result to one of PCA result with entropy value of 0.69 in Figure 3 and one of ARHP result with entropy value of 0.79 in Figure 4, we can see the big differences in the quality of the clusters obtained from these experiments.

Our experiments suggest that clustering methods based on partitioning seem to work best for this type of information retrieval applications, because (1) they do not depend in a choice of a distance function; (2) they do not require calculation of the mean of the clusters, and so the issue of having cluster means very close in space does not apply; (3) they are not sensitive to the dimensionality of the data sets; and (4) they are linearly scalable w.r.t. the cardinalities of the document and feature spaces (in contrast to HAC and AutoClass which are quadratic). In particular, both the hypergraph partitioning method and the principal component methods performed much better than the traditional methods regardless of the feature selection criteria used.

There were also dramatic differences in run times of the four methods. For example,






Figure 3: Class distribution of *PCA* clusters.




Figure 4: Class distribution of *ARHP* clusters.

when no feature selection criteria was used (dataset size of  $185 \times 10538$ ), ARHP and PCA took less than 2 minutes, whereas HAC took 1 hour and 40 minutes and AutoClass took 38 minutes.




Figure 5: Entropy of different algorithms. Note that lower entropy indicates better cohesiveness of clusters.

Aside from overall performance and the quality of clusters, the experiments point to a few other notable conclusions. As might be expected, in general clustering algorithms yield better quality clusters when the full set of feature is used (experiment  $E_1$ ). Of course, as the above discussion shows, for large datasets the computational costs may be prohibitive, especially in the case of HAC and AutoClass methods. It is therefore important to select a smaller set of representative features to improve the performance of clustering algorithms without losing too much quality. Our experiments with various feature selection methods represented in  $E_1$  through  $E_{10}$ , clearly show that restricting the feature set to those only appearing in the frequent item sets (discovered as part of the association rule algorithm), has succeeded in identifying a small set of features that are relevant to the clustering task. In fact, in the case of AutoClass and HAC, the experiment  $E_{10}$  produced results that were better than those obtained by using the full set.

It should be noted that the conclusions drawn in the above discussion have been confirmed by another experiment using a totally independent set of documents [MHB<sup>+</sup>97].

## 5 Search for and Categorization of Similar Documents

One of the main tasks of the agent is to search the Web for documents that are related to the clusters of documents. The key question here is how to find a representative set of words that can be used in a Web search. With a single document, the words appearing in the document become a representative set. However, this set of words cannot be used directly in a search

because it excessively restricts the set of documents to be searched. The logical choice for relaxing the search criteria is to select words that are very frequent in the document.

The characteristic words of a cluster of documents are the ones that have high document frequency and high average text frequency. Document frequency of a word refers to the frequency of the word across documents. Text frequency of a word refers to word frequency within a document. We define the TF word list as the list of  $k$  words that have the highest average text frequency and the DF word list as the list of  $k$  words that have the highest document frequency.

For each cluster, the word lists TF and DF are constructed.  $TF \cap DF$  represents the characteristic set of words for the cluster, as it has the words that are frequent across the document and have high average frequency. The query can be formed as

$$(c_1 \wedge c_2 \dots \wedge c_m) \wedge (t_1 \vee t_2 \dots \vee t_n)$$

where  $c_i \in TF \cap DF$  and  $t_i \in TF - DF$ .

We formed queries from the business capital cluster discussed in Section 4.3. We found the characteristic words of the cluster ( $TF \cap DF$ ) and issued the following query to Yahoo web search engine:

```
+capit* +busi* +financ* +provid* +fund* +develop* +compani* +financi* +manag*
```

The search returned 2280 business related documents. We then added the most frequent words that were not in the previous list ( $TF - DF$ ) to form the following query:

```
+capit* +busi* +financ* +provid* +fund* +develop* +compani* +financi* +manag*
loan* invest* program* credit* industri* tax* increas* cost* technologi*
sba* project*
```

AltaVista search using this query returned only 372 business related documents which seemed highly related to the existing documents in the cluster. First page returned by the query is shown in Figure 6.

The documents returned as the result of queries can be handled in several ways as shown in Figure 1. ARHP could be used to filter out non-relevant documents among the set of documented returned by the query as discussed in Section 4.1. The degree of filtering can be increased either by setting higher support criteria for association rules discovery or by having a tighter connectivity constraint in the partition.

Resulting documents can be incrementally added to the existing clusters using ARHP or PCA depending on the method used for clustering. With ARHP, for each new document, existing hyperedges are extended to include the new document and their weights are calculated. For each cluster, the connectivity of this new document to the cluster is measured by adding the weights of all the extended hyperedges within the cluster. The new document is placed into the cluster with the highest connectivity. The connectivity ratio between the chosen cluster and the remaining clusters indicates whether the new document strongly belongs to the chosen cluster. If the connectivity of the document is below some threshold for all clusters, then the document can be considered as not belonging to any of the cluster.

With PCA, the binary tree can also be used to filter new incoming documents by placing the document on one or the other side of the root hyperplane, then placing it on one or the



[MLB Playoffs](#) - [NHL Preseason](#) -- [Drop-Off Locator](#)

---

Categories - Sites - [Alta Vista Web Pages](#) | [Net Events](#) - [Headlines](#) - [Amazon.com Related Books](#)

---

### Alta Vista Web Pages (1-20 of 372)

- [SBA Loans Take A New Direction](#) - SBA Loans Take A New Direction. April, 1993. While the restrictive conditions in the commercial lending environment show some signs of abating, obtaining..  
--[http://www.ffgroup.com/contractor/sba\\_8a/504loans.html](http://www.ffgroup.com/contractor/sba_8a/504loans.html)
- [SBA: Small Business Act of 1958 and PL 104-208, Approved 9/30/96](#) - This compilation includes PL 104-208, approved 9/30/96. SMALL BUSINESS INVESTMENT ACT OF 1958. (Public Law 85-699, as amended) Sec. 101. SHORT TITLE This..  
--<http://www.sbaonline.sba.gov/INV/sbaact.html>
- [New Haven Enterprise Community Summary](#) - EC Summary Contact EC Summary Maps. STRATEGIC PLAN SUMMARY. Introduction. The New Haven Enterprise Community Strategic Plan marshals our community's..  
--<http://www.hud.gov/cpd/ezec/ct/ctnewhav.html>
- [ABIOTENESIS SOFTWARE - Business Venture Finance Investment Info](#) - The Abiogenesis Business Finance Resource Site. Abiogenesis provides software for the creation of computer dictionaries. Setting up and capitalizing your..  
--<http://www.abiogenesis.com/AbioDocs/Finance.html>
- [Financial Information](#) - Finance Executives. General Web Resources. CorpFiNet. An Introduction to the WWW for Executives. SuperCFOs. Well-written article from Fortune discusses..  
--<http://www.unf.edu/students2/jroger2/finance.html>
- [Fairfax County Business Services and Resources \(Part 3\)](#) - Business Services and Resources. Arts. Associations. Career Development/Continuing Education. Child Care. Chambers of Commerce and Other Business..  
--[http://www.eda.co.fairfax.va.us/fceda/do\\_bus/b\\_resrc3.html\\_\\_28506-4](http://www.eda.co.fairfax.va.us/fceda/do_bus/b_resrc3.html__28506-4)
- [Canadian Financial Regulation: A System in Transition](#) - Commentary 78; Financial Regulation March 19, 1996. Canadian Financial Regulation: A System in Transition. by Edwin H. Neave. Abstract. Planned revisions..  
--<http://www.cdhowe.org/eng/word/word-5.html>
- [FBS | Business Page | re:BUSINESS | Summer 1996](#) - RE: Business. SUMMER 1996. THE FIRST AMERICAN 401(K) SOLUTION FOR EMPLOYEES' RETIREMENT. Today, many businesses are setting up 401(k) plans for..  
--[http://www.fbs.com/biz\\_pages/newsletters/96summer.html](http://www.fbs.com/biz_pages/newsletters/96summer.html)
- [CPB TV Future Fund Business Plans](#) - TV Future Fund. Business Plan Outline. Updated

Figure 6: Search results from Yahoo.

other side the next appropriate hyperplane, letting it percolate down the tree until it reaches a leaf node. This identifies the cluster in the original tree most closely related to the new incoming document. If the combined scatter value for that cluster with the new document is above a given threshold, then the new document is only loosely related to that cluster, which can then be split in two.

## 6 Conclusion

In this paper we have proposed an agent to explore the Web, categorizing the results and then using those automatically generated categories to further explore the Web. We have presented sample performance results for the categorization (clustering) component, and given some examples showing how those categories are used to return to the Web for further exploration.

For the categorization component, our experiments have shown that the ARHP algorithm and the PCA algorithm are capable of extracting higher quality clusters while operating much faster compared to more classical algorithms such as HAC or AutoClass. This is consistent with our previous results [MHB<sup>+</sup>97]. The ARHP algorithm is also capable of filtering out documents by setting a support threshold.

To search for similar documents keyword queries are formed by extending the characteristic word sets for each cluster. Our experiments show that this method is capable of producing small sets of relevant documents using standard search engines.

In the future, we will explore the performance of the entire agent as an integrated and fully automated system, comparing the relative merits of the various algorithms for clustering, query generation, and document filtering, when used as the key components for this agent.

## References

- [Ack97] M. Ackerman et al. Learning probabilistic user profiles. *AI Magazine*, 18(2):47–56, 1997.
- [AFJM95] R. Armstrong, D. Freitag, T. Joachims, and T. Mitchell. WebWatcher: A learning apprentice for the world wide web. In *Proc. AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments*. AAAI Press, 1995.
- [AMS<sup>+</sup>96] A. Agrawal, H. Mannila, R. Srikant, H. Toivonen, and A.I. Verkamo. Fast discovery of association rules. In U.M. Fayyad, G. Piatetsky-Shapiro, P. Smith, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 307–328. AAAI/MIT Press, 1996.
- [BDO95] M. W. Berry, S. T. Dumais, and Gavin W. O’Brien. Using linear algebra for intelligent information retrieval. *SIAM Review*, 37:573–595, 1995.
- [Ber76] C. Berge. *Graphs and Hypergraphs*. American Elsevier, 1976.
- [BSY95] Marko Balabanovic, Yoav Shoham, and Yeogirl Yun. An adaptive agent for automated Web browsing. *Journal of Visual Communication and Image Representation*, 6(4), 1995.
- [CS96] P. Cheeseman and J. Stutz. Bayesian classification (autoclass): Theory and results. In U.M. Fayyad, G. Piatetsky-Shapiro, P. Smith, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 153–180. AAAI/MIT Press, 1996.
- [DEW96] R. B. Doorenbos, O. Etzioni, and D. S. Weld. A scalable comparison shopping agent for the World Wide Web. Technical Report 96-01-03, University of Washington, Dept. of Computer Science and Engineering, 1996.

- [DH73] Richard O. Duda and Peter E. Hart. *Pattern Classification and scene analysis*. John Wiley & Sons, 1973.
- [FBY92] W. B. Frakes and R. Baeza-Yates. *Information Retrieval Data Structures and Algorithms*. Prentice Hall, Englewood Cliffs, NJ, 1992.
- [Fra92] W. B. Frakes. Stemming algorithms. In W. B. Frakes and R. Baeza-Yates, editors, *Information Retrieval Data Structures and Algorithms*, pages 131–160. Prentice Hall, 1992.
- [GV96] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins Univ. Press, 3rd edition, 1996.
- [HBML95] K. Hammond, R. Burke, C. Martin, and S. Lytinen. FAQ-Finder: A case-based approach to knowledge navigation. In *Working Notes of the AAAI Spring Symposium: Information Gathering from Heterogeneous, Distributed Environments*. AAAI Press, 1995.
- [HKKM97] E.H. Han, G. Karypis, V. Kumar, and B. Mobasher. Clustering based on association rule hypergraphs. In *Workshop on Research Issues on Data Mining and Knowledge Discovery*, pages 9–13, Tucson, Arizona, 1997. An extended version appears in TR-97-019, Dept of Computer Science and Engineering, University of Minnesota, Minneapolis.
- [JD88] A.K. Jain and R. C. Dubes. *Algorithms for Clustering Data*. Prentice Hall, 1988.
- [KAKS97] G. Karypis, R. Aggarwal, V. Kumar, and S. Shekhar. Multilevel hypergraph partitioning: Application in VLSI domain. In *Proceedings ACM/IEEE Design Automation Conference*, 1997.
- [KLSS95] T. Kirk, A. Y. Levy, Y. Sagiv, and D. Srivastava. The information manifold. In *Working Notes of the AAAI Spring Symposium: Information Gathering from Heterogeneous, Distributed Environments*. AAAI Press, 1995.
- [KW96] C. Kwok and D. Weld. Planning to gather information. In *Proc. 14th National Conference on AI*, 1996.
- [LF78] S.Y. Lu and K.S. Fu. A sentence-to-sentence clustering procedure for pattern analysis. *IEEE Transactions on Systems, Man and Cybernetics*, 8:381–389, 1978.
- [MHB<sup>+</sup>97] J. Moore, S. Han, D. Boley, M. Gini, R. Gross, K. Hastings, G. Karypis, V. Kumar, and B. Mobasher. Web page categorization and feature selection using association rule and principal component clustering. In *7th Workshop on Information Technologies and Systems*, Dec 1997.
- [MS96] Y. S. Maarek and I.Z. Ben Shaul. Automatically organizing bookmarks per content. In *Proc. of 5th International World Wide Web Conference*, 1996.
- [PE95] M. Perkowitz and O. Etzioni. Category translation: learning to understand information on the internet. In *Proc. 15th International Joint Conference on AI*, pages 930–936, Montral, Canada, 1995.
- [Por80] M. F. Porter. An algorithm for suffix stripping. *Program*, 14(3):130–137, 1980.
- [SM83] Gerard Salton and Michael J. McGill. *Introduction to Modern Information Retrieval*. McGraw-Hill, 1983.
- [TSM85] D.M. Titterington, A.F.M. Smith, and U.E. Makov. *Statistical Analysis of Finite Mixture Distributions*. John Wiley & Sons, 1985.
- [WP97] Marilyn R. Wulfekuhler and William F. Punch. Finding salient features for personal Web page categories. In *Proc. of 6th International World Wide Web Conference*, April 1997.
- [WVS<sup>+</sup>96] Ron Weiss, Bienvenido Velez, Mark A. Sheldon, Chanathip Nemprempre, Peter Szilagyi, Andrzej Duda, and David K. Gifford. Hypursuit: A hierarchical network search engine that exploits content-link hypertext clustering. In *Seventh ACM Conference on Hypertext*, March 1996.