

Data Mining Approaches to User Modeling for Adaptive Hypermedia: Survey and Future Directions

Enrique Frias-Martinez, Sherry Y. Chen and Xiaohui Liu

Abstract—The ability of an adaptive hypermedia system to create tailored environments depends mainly on the amount and accuracy of information stored in each user model. Some of the difficulties that user modeling faces are the amount of data available to create user models, the adequacy of the data, the noise within that data and the necessity of capturing the imprecise nature of human behavior. Data mining and machine learning techniques have the ability to handle large amounts of data and to process uncertainty. These characteristics make these techniques suitable for automatic generation of user models that simulate human decision-making. This paper surveys different data mining techniques that can be used to efficiently and accurately capture user behavior. The paper also presents guidelines that show which techniques may be used more efficiently according to the task implemented by the application.

Index Terms— Adaptive Hypermedia, Data Mining, Machine Learning, User Modeling.

I. INTRODUCTION

Adaptive hypermedia can be defined as the technology that allows to personalize for each individual user of a hypermedia application the content and presentation of the application according to user preferences and characteristics (Perkowitz et al., 1998) (Perkowitz et al., 1999) (Perkowitz et al., 2000).

The process of personalization is defined as the ways in which information and services can be tailored to match the unique and specific needs of an individual or a community (Callan et al., 2001). Personalization is about building customer loyalty by developing a meaningful one-to-one relationship; by understanding the needs of each individual and helping satisfy a goal that efficiently and knowledgeably addresses each individual's need in a given context (Riecken, 2000). The more information a user model has, the better the content and presentation will be personalized. A user model is created through a user modeling process in which unobservable information about a user is inferred from observable information from that user; for example, using the interactions with the system (Zukerman, Albrecht & Nicholson, 1999).

This work was supported in part by the UK Arts and Humanities Research Board (AHRB) under Grant MRG/AN9183/APN16300.

E. Frias-Martinez, S. Chen and X. Liu are with Department of Information Systems & Computing, Brunel University, Uxbridge, Middlesex UB8 3PH, United Kingdom ; phone: +44 (0)1895 266023; fax: +44 (0) 1895251686; e-mail: {enrique.frias-martinez, sherry.chen, xiaohui.liu}@brunel.ac.uk.

User models can be created using a user-guided approach, in which the models are directly created using the information provided by each user, or an automatic approach, in which the process of creating a user model is controlled by the system and is hidden from the user. The user-guided approach produces adaptable services and adaptable user-models (Fink et al., 1997), while the automatic approach produces adaptive services and adaptive user models (Fink et al., 1997) (Brusilovsky et al., 1997). In general, a user model will contain some adaptive and some adaptable elements. Ideally the set of adaptable elements should be reduced to the minimum possible (elements like age, gender, favorite background color, etc.), while other elements (favorite topics, patterns of behavior, etc.) should be created by a learning process. These concepts have also been presented in the literature as implicit and explicit user modeling acquisition (Quiroga et al., 1999).

The problem of user modeling can be implemented using an automatic approach because a typical user exhibits patterns when accessing a hypermedia system and the set of interactions containing those patterns can be stored in a log database in the server. In this context, machine learning and data mining techniques can be applied to recognize regularities in user trails and to integrate them as part of the user model. Machine learning encompasses techniques where a machine acquires/learns knowledge from its previous experience (Witten & Frank, 1999). The output of a machine learning technique is a structural description of what has been learned that can be used to explain the original data and to make predictions. From this perspective, data mining and other machine learning techniques make it possible to automatically create user models for the implementation of adaptive hypermedia services. These techniques make it possible to create user models in an environment, such as a hypermedia application, in which, usually, users are not willing to give feedback of their actions.

This fact leads researchers to consider the different techniques that can be used to capture and model user behavior and which elements of the behavior of a user each one of those techniques can capture. In general the adaptive service that is going to be implemented determines the information that a user model should contain. Once that is known, and taking into account the characteristics of the available data, different data mining/machine learning

techniques can be chosen to capture the necessary patterns.

This paper has surveyed the development of user models using data mining and machine learning techniques from 1999 to 2004, focusing on the main journals and conferences for user modeling, mainly: User Modeling and User-Adapted Interaction, International Conference on User Modeling, IEEE Transactions on Neural Networks, Workshop of Intelligent Techniques for Web Personalization (part of IJCAI International Joint Conference of Artificial Intelligence) and International Workshop on Knowledge Discovery on the WEB (WEBKDD, part of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining). The paper's intentions are (1) to give an up-to-date view of data mining techniques applied to User Modeling and highlight their potential advantages and limitations, and (2) to give basic guidelines about which techniques can be useful for a given adaptive application. The paper complements the results of (Zuckerman et al., 2001) that reviewed the use of predictive statistical models for user modeling.

The organization of the paper is as follows. The paper first defines the concept of user model, its relevance for adaptive hypermedia and the basic steps for the automatic creation of user models. After that a set of unsupervised and supervised techniques are presented. For each technique we present its theoretical background, its pros and cons and its applications in the field of user modeling. Next we develop a guideline on how to create a user model according to the needs of the adaptive hypermedia application that is going to be implemented. The conclusions section closes the paper.

II. USER MODELING (UM)

A user model should capture the behavior (patterns, goals, interesting topics, etc.) a user shows when interacting with the Web. A user model is defined as a set of information structures designed to represent one or more of the following elements (Kobsa, 2001): (1) representation of goals, plans, preferences, tasks and/or abilities about one or more types of users; (2) representation of relevant common characteristics of users pertaining to specific user subgroups or stereotypes; (3) the classification of a user in one or more of these subgroups or stereotypes; (4) the recording of user behavior; (5) the formation of assumptions about the user based on the interaction history and/or (6) the generalization of the interaction histories of many users into groups.

A. User Modeling and Adaptive Hypermedia

Adaptive hypermedia allows to personalize to each individual the content and presentation of a web site. The architecture of an adaptive hypermedia system is usually divided in two parts: the server side and the client side. The server side generates the user models from a data base containing the interactions of the users with the system and the personal data/preferences that each user has given to the system. Also, user models can contain knowledge introduced by the designer (in the form of rules for example). These user models, in combination with a hypermedia database, are used by the "Decision Making and Personalization Engine" module

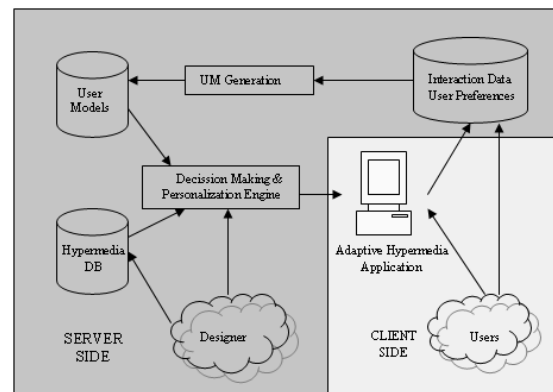


Fig. 1. Generic Architecture of an Adaptive Hypermedia Application.

to identify user needs, decide on the types of adaptation to be performed and communicate them to an adaptive interface. Fig. 1 presents the architecture of a generic AH system.

A personalized hypermedia system, by its very nature, should respond in real time to user inputs. To do so, the architecture of the system should provide a quick access to the right information at the right time.

Adaptive hypermedia uses the knowledge given by user models to implement one or more of the two basic types of adaptive tasks:

- *Recommendation (R)*. Recommendation is the capability of suggesting interesting elements to a user based on some information; for example from the items to be recommended or from the behavior of other users. Recommendation is also known in the literature as collaborative filtering.

- *Classification (C)*. Classification builds a model that maps or classifies data items into one of several predefined classes. Classification is done using only data related to that particular item. This knowledge can be used to tailor the services of each user stereotype (class). In the literature, classification has also been presented as content-based filtering.

Both recommendation and classification are types of filtering applications.

B. Automatic Generation of User Models

One of the processes presented in Fig. 1 is the automatic generation of user models from the interaction data between the users and the system (done by the UM Generation module). Fig. 2 presents the basic steps of this module: (1) Data Collection, (2) Preprocessing, (3) Pattern Discovery and (4) Validation and Interpretation:

- *Data Collection*. In this stage user data is gathered. For automatic user modelling the data collected includes: data regarding the interaction between the user and the Web, data regarding the environment of the user when interacting with the Web, direct feedback given by the user, etc.

- *Data Preprocessing/Information Extraction*. The information obtained in the previous stage cannot be directly processed. It needs to be cleaned from noise and inconsistencies in order to be used as the input of the next

phase. For user modelling, this involves mainly user identification and session reconstruction. This stage is aimed at obtaining, from the data available, the semantic content about the user interaction with the system. Also in this phase the data extracted should be adapted to the data structure used by standard pattern discovery algorithms used in the next step.

- **Pattern Discovery.** In this phase, machine learning and data mining techniques are applied to the data obtained in the previous stage in order to capture user behaviour. The output of this stage is a set of structural descriptions of what have been learned about user behaviour and user interests. These descriptions constitute the base of a user model. Different techniques will capture different user properties and will express them in different ways. The knowledge needed to implement an adaptive service will determine, among other factors, which techniques to apply in this phase.

- **Validation and Interpretation.** In this phase the structures obtained in the pattern discovery stage are analyzed and interpreted. The patterns discovered can be interpreted and validated, using domain knowledge and visualization tools, in order to test the importance and usability of the knowledge obtained. In general this process is done with the help of a user modelling designer.

User model heuristics are used in each step of the process to extract, adapt and present the knowledge in a relevant way. The process of generation of user models using data mining/machine learning techniques can be seen as a standard process of extracting knowledge from data where user modelling is used as a wrapper for the entire process.

C. Data Mining and its Relevance to User Modeling

As it has been presented in Fig. 2, the phase of Pattern Discovery finds out relevant information about the behavior of a user (or set of users) when interacting with the Web. Data mining and machine learning techniques are ideal for that process because they are designed to represent what have been learned from the input data with a structural representation. This representation stores the knowledge needed to implement the two types of tasks previously described.

Although the application of machine learning and data mining techniques work really well for modeling user behavior, it also faces some problems. As stated in (Webb et al., 2001), among those challenges is the problem of needing large data sets, the problem of labeling data for supervised machine learning techniques and the problem of computational complexity.

Each data mining/machine learning technique will capture different relationships among the data available and will express the results using different data structures. The key question is to find out which patterns need to be captured in order to implement an adaptive service. It is important, in order to choose a suitable learning method, to know what knowledge is captured by each technique and how that knowledge can be used to implement the two basic tasks. Also, the choice of learning method depends largely on the type of training data available. The main distinction in




Fig. 2. Steps for Automatic Generation of User Models.

machine learning research is between supervised and unsupervised learning.

Supervised learning requires the training data to be preclassified. This means that each training item is assigned a unique label, signifying the class to which the item belongs. Given these data, the learning algorithm builds a characteristic description for each class, covering the examples of this class. The important feature of this approach is that the class descriptions are built conditional to the preclassification of the examples in the training set.

In contrast, unsupervised learning methods do not require preclassification of the training examples. These methods form clusters of examples, which share common characteristics. The main difference to supervised learning is that categories are not known in advance, but constructed by the learner itself. When the cohesion of a cluster is high, i.e., the examples in it are similar, it defines a new class.

The rest of the paper presents how data mining and machine learning techniques have been used for user modeling: which knowledge can be captured with each technique, examples of applications and its limits and strengths. The techniques presented are divided into two groups: unsupervised, which includes clustering (hierarchical, non-hierarchical and fuzzy clustering) and association rules, and supervised, which includes Decision trees/Classification rules, k-nearest neighbor (k-NN), Neural Networks and Support Vector Machines (SVM).

III. UNSUPERVISED APPROACHES TO USER MODELING

The main unsupervised techniques are clustering and association rules. Clustering comprises a wide variety of different techniques based in the same concept. A collection of different clustering techniques and its variations can be found in (Jain et al., 1988).

A. Clustering for User Modeling

The task of clustering is to structure a given set of unclassified instances (data vectors) by creating concepts, based on similarities found on the training data.

A clustering algorithm finds the set of concepts that cover all examples verifying that: (1) the similarity between examples of the same concepts is maximized, and (2) the similarity between examples of different concepts is

minimized. In a cluster algorithm the key element is how to obtain the similarity between two items of the training set.

Clustering techniques can be classified in hard clustering and fuzzy clustering. In non-fuzzy or hard clustering, data is divided into crisp clusters, where each data point belongs to exactly one cluster. In fuzzy clustering, the data points can belong to more than one cluster, and associated with each of the instances are membership grades which indicate the degree to which they belong to the different clusters.

Hard clustering techniques may be grouped into two categories: hierarchical and non-hierarchical (Jain, 1999). A hierarchical clustering procedure involves the construction of a hierarchy or tree-like structure, which is basically a nested sequence of partitions, while non-hierarchical or partitional procedures end up with a particular number of clusters at a single step.

1) Basic Algorithms: Non-hierarchical Techniques

The main non-hierarchical clustering techniques are: (1) k-means clustering and (2) Self-Organizing Maps (SOM).

a) K-means Clustering

The k-means clustering technique (MacQueen, 1967) is given as an input the number of clusters k . The algorithm then picks k items, called seeds, from the training set in an arbitrary way. Then, in each iteration, each input item is assigned to the most similar seed, and the seed of each cluster is recalculated to be the centroid of all items assigned to that seed. This process is repeated until the seed coordinates stabilize. This algorithm aims at minimizing an objective function, J , typically a squared error function:

$$J = \sum_{j=1}^k \sum_{i=1}^n d_{ij} = \sum_{j=1}^k \sum_{i=1}^n \|x_i^{(j)} - c_j\|^2 \quad (1)$$

where d_{ij} is the distance measure between a data point x_i and the cluster centre c_j . J is an indicator of the distance of the n data points from their respective cluster centres and it represents the compactness of the clusters created.

Although it can be proved that the procedure will always terminate, the k-means algorithm does not necessarily find the most optimal configuration, corresponding to the global objective function minimum. The K-means algorithm is popular because it is easy to understand and easy to implement. Its main drawback is that it can take a considerable time if the number of patterns involved is very large and the number of clusters is substantial. Another problem is that it is sensitive to the initial partition – the selection of the initial patterns, and may converge to a *local minimum* of the criterion function value if the initial partition is not properly chosen. A possible remedy is to run the algorithm with a number of different initial partitions. If they all lead to the same final partition, this implies that the global minimum of the square error has been achieved. However, this can be time-consuming, and may not always work.

b) Self Organizing Maps (SOM)

The self-organising map (SOM) algorithm was proposed by Teuvo Kohonen in 1981 (Kohonen, 1997). Apart from being used in a wide variety of fields, the SOM offers an interesting tool for exploratory data analysis, particularly for partitional clustering and visualization. It is capable of representing high-dimensional data in a low dimensional space (often a two or one dimensional array) that preserves the structure of the original data.

A self-organising network consists of a set of input nodes $V = \{v_1, v_2, \dots, v_N\}$, a set of output nodes $C = \{c_1, c_2, \dots, c_M\}$, a set of weight parameters $W = \{w_{11}, w_{12}, \dots, w_{ij}, \dots, w_{NM}\}$ ($1 \leq i \leq N$, $1 \leq j \leq M$, $0 \leq w_{ij} \leq 1$), and a map topology that defines the distances between any given two output nodes. The output nodes in SOM are usually arranged in a 2-dimensional array to form a “map”. Each input node is fully connected to every output node via a variable connection. A weight parameter is associated with each of these connections, and the weights between the input nodes and output nodes are iteratively changed during the learning phase until a termination criterion is satisfied. For each input vector v , there is one associated *winner node* on the output map. A winner node is an output node that has minimum distance to the input vector.

The SOM algorithm starts by initializing the topology and size of the output map (M), the connection weights to random values over the interval $[0, 1]$, the gain value η (learning rate) and the neighbourhood size r , and normalizes both the input vectors and the connected weight vectors. After that, the algorithm calculates the Euclidean distance between the new input vector v and each node on the output map, and designates the node with the minimum distance as the winner node c :

$$c = \min \sqrt{\sum_{k=1}^N (v_k - w_{kj})^2}, j=1,2,\dots,M \quad (2)$$

Once c is obtained, the algorithm updates the weights W , the learning rate η and N_c , the neighbourhood surrounding the winner node c , in such way that the vectors represented by output nodes are similar if they are located in a small neighborhood. For each node $j \in N_c$, the SOM algorithm performs the following operation:

$$w_j^{(new)} = w_j^{(old)} + \eta [v_i - w_j^{(old)}] \quad (3)$$

After that the neighborhood size and the learning rate are decreased. This process is repeated until it converges. The neighborhood set N_c is a set of nodes that surround the winner node c . These nodes in N_c are affected with weight modifications, apart from those changes made to the winner node, as defined in the algorithm. These weight changes are made to increase the matching between the nodes in N_c and the corresponding input vectors. As the update of weight parameters proceeds, the size of the neighborhood set is slowly decreased to a predefined limit, for instance, a single node. This process leads to one of the most important

properties of SOM that similar input vectors are mapped to geometrically close winner nodes on the output map. This is called *neighborhood preservations*, which has turned out to be very useful for clustering similar data patterns.

It is not always straightforward to visually inspect the projected data on the two-dimensional output map in order to decide the number and size of natural clusters. Therefore, careful analysis or post-processing of output maps is crucial to the partition of the original data set. Like the K-means algorithm, the SOM produces a sub-optimal partition if the initial weights are not chosen properly. Moreover, its convergence is controlled by various parameters such as the learning rate, the size and shape of the neighbourhood in which learning takes place. Consequently, the algorithm may not be very *stable* in that a particular input pattern may produce different winner nodes on the output map at different iterations.

2) Basic Algorithms: Hierarchical Techniques

The main problem of non-hierarchical approaches is that when working with high dimensional problems, in general, there will not be enough items to populate the vector space, which will imply that most dimensions will be unreliable for similarity computations. In order to solve this problem hierarchical clustering techniques were developed. There are two types of hierarchical clustering: agglomerative and divisive. Both share a common characteristic: they create a hierarchy of clusters. While the agglomerative approach creates a bottom-up hierarchy the divisive approach creates a top-down one. Generally speaking, divisive algorithms are computationally less efficient.

A typical hierarchical agglomerative clustering algorithm is outlined below:

- 1) Place each pattern in a separate cluster
- 2) Compute the proximity matrix of all the inter-pattern distances for all distinct pairs of patterns
- 3) Find the most similar pair of clusters using the matrix. Merge these two clusters into one, decrement number of clusters by one and update the proximity matrix to reflect this merge operation.
- 4) If all patterns are in one cluster, stop. Otherwise, go to the above step 2.

The output of such algorithm is a nested hierarchy of trees that can be cut at a desired dissimilarity level forming a partition. Hierarchical agglomerative clustering algorithms differ primarily in the way they measure the distance or similarity of two clusters where a cluster may consist of only a single object at a time. The most commonly used inter-cluster measures are:

$$d_{AB} = \min_{\substack{i \in A \\ j \in B}}(d_{ij}), \quad (4)$$

$$d_{AB} = \max_{\substack{i \in A \\ j \in B}}(d_{ij}), \quad (5)$$

$$d_{AB} = \frac{1}{n_A n_B} \sum_{i \in A} \sum_{j \in B} d_{ij}, \quad (6)$$

where d_{AB} is the dissimilarity between two clusters A and B , d_{ij} is the dissimilarity between two individual patterns i and j , n_A and n_B are the number of individuals in clusters A and B respectively. These three inter-cluster dissimilarity measures are the basis of the three of the most popular hierarchical clustering algorithms. The *single-linkage* algorithm uses (4), the *minimum* of the distances between all pairs of patterns drawn from the two clusters (one pattern from each cluster). The *complete-linkage* algorithm uses Equation (5), the *maximum* of all pairwise distances between patterns in the two clusters. The *group-average* algorithm uses Equation (6), the average of the distances between all pairs of individuals that are made up of one individual from each cluster.

A challenging issue with hierarchical clustering is how to decide the *optimal* partition from the hierarchy. One approach is to select a partition that best fits the data in some sense, and there are many methods that have been suggested in the literature (Everitt, 1993). It has also been found that the single-linkage algorithm tends to exhibit the so-called *chaining* effect: it has a tendency to cluster together at a relatively low level objects linked by chains of intermediates. As such, the method is appropriate if one is looking for “optimally” connected clusters rather than for homogeneous spherical clusters. The complete-linkage algorithm, on the other hand, tends to produce clusters that tightly bound or compact, and has been found to produce more useful hierarchies in many applications than the single-link algorithm (Jain, 1999). The group-average algorithm is also widely-used. Detailed discussion and practical examples of how these algorithms work can be found in (Jain, 1999) and (Webb, 1999).

3) Basic Algorithms: Fuzzy Clustering

One of the most widely used fuzzy clustering algorithms is the Fuzzy C-Means (FCM) Algorithm (Bezdek, 1981). The FCM algorithm attempts to partition a finite collection of elements $X = \{x_1, \dots, x_n\}$ into a collection of c fuzzy clusters with respect to some given criterion. Given a finite set of data, the algorithm returns a list of c cluster centres $C = \{c_1, \dots, c_c\}$ and a partition matrix $U = u_{ij} \in [0, 1]$, $i = 1, \dots, n$, $j = 1, \dots, c$, where each element tells the degree to which element x_i belongs to cluster c_j . Like the k-means algorithm, the fuzzy c-means aims to minimise an objective function. The standard function is:

$$J = \sum_{j=1}^c \sum_{i=1}^n (u_{ij})^m \|x_i^{(j)} - c_j\|^2 \quad (7)$$

which differs from the k-means objective function by the addition of the membership values u_{ij} and the fuzzifier m . The

fuzzifier m determines the level of cluster fuzziness. A large m results in smaller memberships u_{ij} and hence, fuzzier clusters. In the limit $m=1$, the memberships u_{ij} converge to 0 or 1, which implies a crisp partitioning. In the absence of experimentation or domain knowledge, m is commonly set to 2. The basic Fuzzy C-Means Algorithm, given n data points (x_1, \dots, x_n) to be clustered, a number of c clusters with (c_1, \dots, c_c) the center of the clusters, and m the level of cluster fuzziness with,

$$m \in \mathbb{R} > 1, \quad (8)$$

first initializes the membership matrix U to random values, verifying that:

$$u_{ij} \in [0,1], \sum_{j=1}^c u_{ij} = 1 \quad (9)$$

After the initialization, the algorithm obtains the center of the clusters $c_j, j=1, \dots, c$:

$$c_j = \frac{\sum_{i=1}^n (u_{ij})^m x_i}{\sum_{i=1}^n (u_{ij})^m} \quad (10)$$

And obtains the distance between all points $i=1, \dots, n$ and all cluster centers $j=1, \dots, c$

$$d_{ij} = \|x_i^{(j)} - c_j\| \quad (11)$$

Updating matrix U according to the new distances,

$$d_{ij} = 0 \Rightarrow u_{ij} = 1$$

$$u_{ij} = \left[\sum_{k=1}^c \left(\frac{d_{ij}}{d_{ik}} \right)^{\frac{2}{m-1}} \right]^{-1} \quad (12)$$

This process is repeated until the set of cluster centers is stabilized. There are other algorithms, which are optimizations of the original FCM, like Fuzzy c-Medoid Algorithm (FCMdd) or the Fuzzy c-Trimmed Medoids Algorithm (FCTMdd) (Krishnapuram et al., 2001)

4) Applications for User Modeling

For UM there are two kinds of interesting clusters to be discovered: usage clusters and page clusters. Clustering of users tends to establish groups of users exhibiting similar browsing patterns, which are usually called stereotypes. Such knowledge is especially useful for inferring user demographics in order to perform market segmentation in e-commerce-applications or provide personalized Web content to the users. On the other hand, clustering of pages will discover groups of pages having related content. This information is useful for Internet search engines and Web assistance providers.

In the context of UM, clustering has a distinguishable characteristic: it is usually done with non-numerical data. This implies that, usually, the clustering techniques applied are relational, where numerical values represent the degrees to which two objects of the data set are related. Clustering applied to user modeling has to use techniques that can handle relational data because the information used to create clusters (pages visited, characteristics of the user, etc.) cannot usually be represented by numerical vectors. In case they are represented using vectors, part of the semantic of the original data is lost. In these systems the definition of distance is done using vectorial representations of user interactions with the adaptive hypermedia system. Some examples of relational

TABLE I
EXAMPLES OF SOME CLUSTERING-BASED USER MODELS

	Application	Input Data	Results
(Mobasher et al, 2000)	Capture of web-users interests using k-means Clustering. Design of a cluster-based recommendation system	User logs from the Univ. of Minnesota Comp. Science web server collected during a month.	Example of the implementation of the recommendation system in a commercial site.
(Paliouras et al. 1999)	News- filtering system based on communities of users. Clustering allows to recommend interesting news to a user.	When registering a user specifies his/her interests.	Established machine learning techniques are very useful for the acquisition of communities of users
(Doux et al., 1997)	K-means clustering algorithm for user profiling in order to derive prototypical behavior from each user.	Data collected from a dedicated set of experiments where users are asked about their preferences.	The techniques proposed handle qualitative data for clustering users efficiently.
(Fu et al., 1999)	Grouping of users with a common behavior in a web server taking into account access patterns .	Data collected from UMR web server log (www.umd.edu) containing 2.5 million records.	The clusters obtained can be used for personalization purposes.
(Hay et al., 2001)	Clustering methods that capture the inherent sequentiality of web visits. A metric, Sequence Alignment Method, is introduced to be used instead of Euclidean distance for clustering purposes.	Log files of a Belgian telecom provider collected over a one-week period.	The results are as good as the ones obtained with Euclidean distance, while keeping the concept of order.
(Mobasher et al., 2001)	Clustering for collaborative filtering	12,000 sessions collected from the Association for Consumer Research web site.	With the proper data preprocessing the clustering approach outperforms more traditional approaches to this problem (like Nearest Neighbor)

TABLE II
EXAMPLES OF SOME FUZZY CLUSTERING-BASED USER MODELS

	Application	Input Data	Outcome
Lampinen and Koivisto (2002)	Obtain application profiles from network traffic data to manage network resources.	274000 samples of different applications from an edge router of a LAN network.	FCM produced better results than SOM. A method for the comparison of both solutions is also introduced.
Nasraoui et al. (1999)	A new algorithm (CARD) to mine user profiles from access logs is proposed.	12 day log data of the Dep. of Comp. Eng., at Univ. of Missouri.	CARD is very effective for clustering many different profiles in user sessions.
Joshi et al. (2000)	Two algorithms to mine user profiles: FCMdd and FCTMdd.	CSEE logs of Univ. of Maryland	Both algorithms extract interesting user profiles. FCMdd is not able to handle noise as effectively as FCTMdd.
Krishnapuram et al. (2001)	Web access log analysis for user profiling using RFCMdd (Robust Fuzzy c-Medoids).	Five days of CSEE web server activity of Univ. of Maryland.	RFMdd is very effective for clustering of relational data.

clustering applied to web mining are (Joshi et al., 2000) and (Mobasher et al., 2000).

Table I summarizes some studies and applications of “hard” clustering for UM. Some examples of recommendation implemented with clustering algorithms are (Mobasher et al, 2000), (Doux et al., 1999), (Paliouras et al., 1999), which uses SOM, (Fu et al., 2001) and (Mobasher et al, 2001). Examples of classification tasks implemented using clustering are (Doux et al. 1997) and (Hay et al., 2001). (Goren-Bar et al., 2001) uses SOM to classify documents based on a subjectively predefined set of clusters in a specific domain.

When using Fuzzy Clustering, the main difference with hard clustering techniques is that, a user can be at the same time in more than one cluster with different degrees of truth. This allows to better capture the inherent uncertainty that the problem of modeling user behavior has. Examples of applications that implement a recommendation task using FC include (Lampinen and Koivisto, 2002), (Nasraoui et al., 1999) and (Nasraoui and Krishnapuram, 2000). Examples of classification tasks are (Joshi et al., 2000) and (Krishnapuram et al., 2001). Table II summarizes some studies and applications of FC for UM.

5) Limitations

The main problem that clustering techniques face is how to define the concept of distance that is going to be used. In general some knowledge of the problem is needed to define an optimum concept of distance. When applied to user modelling this problem is even harder due to the nature of the data available: interactions, user preferences, pages visited, etc., which are not expressed in a numerical way. Different techniques to characterize web user behaviour using numerical vectors have been proposed (Joshi et al., 2000) (Mobasher et al., 2000), but in one way or another, the representations lose part of the semantics that the original data had. More research is needed on how to express the data available in a numerical way and how to define a meaningful concept of distance for creating efficient user models using clustering techniques while at the same time keeping the semantics of the original data.

Non-hierarchical clustering techniques assume that the number of clusters k is known a priori. For user modeling this is not usually the case. This implies that some heuristics need

to be used to determine the number of clusters.

Also user models developed so far using fuzzy clustering do not fully use the fuzzy nature of the technique in order to create more flexible and adaptive systems. More studies are needed to create meaningful ways of mixing the different personalizations associated with each one of the clusters in which a user can be included when using fuzzy clustering.

B. Association Rules for User Modeling

Association rule discovery aims at discovering all frequent patterns among transactions. The problem was originally introduced by (Agrawal et al., 1993) and was based on detecting frequent items in a market basket. The *Apriori* algorithm, including its variations and extensions, is widely accepted to solve this problem.

1) Basic Algorithms

Being $I = \{i_1, i_2, \dots, i_m\}$ a set of items, S a set of transactions, where each transaction T is a set of items, and X a set of items in I , a transaction T is said to contain X , if :

$$X \subseteq T \quad (13)$$

An Association Rule is defined as an implication of the form:

$$X \Rightarrow Y, X \subset I, Y \subset I, X \cap Y = \emptyset \quad (14)$$

Where X is usually called antecedent and Y consequent. The order of the items of both the antecedent and consequent is not relevant. Each Association Rule has a confidence and a support associated. The support of a rule is defined as the fraction of strings in the set of sessions of S where the rule successfully applies:

$$\theta(XY) = \frac{|S_i \in S / XY \in S_i|}{|S|} \quad (15)$$

The confidence of a rule is defined as the fraction of times for which if the antecedent X is satisfied, the consequent Y is also true:

$$\sigma(XY) = \frac{\theta(XY)}{\theta(X)} \quad (16)$$

The values of support and confidence are used to filter the set of association rules obtained from S . Given a set of transactions S , the task of mining association rules can be formulated as finding all association rules with at least minimum support and minimum confidence, where the thresholds for support and confidence are user-specified values. The task of discovering association rules can be done in three steps: (1) find all sets of items that have transaction support above minimum support, (2) for each item obtained, find all non-empty subsets and (3) for each such subset X of item I , if the confidence is bigger than the given threshold, produce the rule:

$$X \Rightarrow (I - X) \quad (17)$$

2) Applications for User Modeling

In the context of user modeling, association rules capture sets of actions that have a causal relation among them. A typical application of association rules for user modeling is capturing pages that are accessed together.

Typically, association rules are used to implement recommendation tasks. Some examples of recommendation tasks implemented using association rules are (Chen et al., 2002), (Lin et al., 2001) and (Nanopoulos et al., 2001), (Mobasher et al., 2001), and (Geyer-Schulz et al., 2002). Table III summarizes these examples.

3) Limitations

Association rules have two important drawbacks when used for user modeling: they do not capture the sequentiality of both the antecedent and the consequent and no temporality information (for example when X happens when is Y going to happen) is captured. Both temporality and sequentiality are very important characteristics for user modeling. Temporality allows us to predict not only what actions a user is going to take, but also when those actions are going to be taken. Sequentiality captures the order in which a user makes a set of actions which is relevant to the consequent that such

antecedent will fire. Because temporality and sequentiality are two key elements that need to be captured to create efficient user models, more research is needed to capture these characteristics when creating association rules.

IV. SUPERVISED APPROACHES TO USER MODELING

This section gives a review of how supervised learning techniques (decision trees/Classification rules, Neural Networks, k-nearest neighbor and Support Vector Machines) can be used to model user behavior.

A. Decision Trees/Classification Rules for User Modeling

Decision tree learning (Mitchell, 1997)(Winston, 1992) is a method for approximating discrete-valued functions with disjunctive expressions. Decision tree learning is generally best suited to problems where instances are represented by attribute-value pairs and the target function has discrete output values.

Classification rules are an alternative representation of the knowledge obtained from classification trees. They construct a profile of items belonging to a particular group according to their common attributes.

1) Basic Algorithms

The training process that creates a decision tree is called induction. A standard decision tree algorithm has two phases: (1) tree growing and (2) pruning. The growing phase can be done using two methods: (1) Top-Down induction and (2) Incremental induction.

Top-down induction is an iterative process which involves splitting the data into progressively smaller subsets. Each iteration considers the data in only one node. The first iteration considers the root node that contains all the data. Subsequent iterations work on derivative nodes that will contain subsets of the data. The algorithm begins by analyzing the data to find the independent variable that, when used as a splitting rule will result in nodes that are most different from each other with respect to the dependent variable. The quality of a test is measured by the impurity/variance of example sets. The most common measure is the information gain. Typically,

TABLE III
EXAMPLES OF SOME ASSOCIATION RULES-BASED USER MODELS

	Application	Input Data	Results
(Chen et al., 2002)	Prediction of user intention in machine-human interaction. The model is used to predict human actions for an Internet browser.	Interaction of 5 users with the system during a month.	The approach achieves 84% average accuracy in predicting user's intention.
(Lin et al., 2001)	Learning user preference models of multimedia Internet files. The model is used to predict when a user will need a multimedia file and to help in the search of the file needed.	200 records of user's emails.	The average correct acceptance and average correct refusal of each user preference modeling is higher than 80%.
(Mobasher et al., 2001)	Efficient personalization of web sites using association rules and non-labeled data contained in server's log file.	Access logs for the Web site Association for Consumer Research (ACR) Newsletter (www.acr-news.org).	Association Rules achieve better recommendation effectiveness than traditional collaborative filtering approaches.
(Geyer-Schulz, 2002)	Evaluation of the recommendations given in a educational Internet information broker using Association Rules and repeat-buying theory.	Data set obtained from the Virtual University Information system of Vienna University during 6 months of usage.	Association Rules tend to work better and achieve higher prediction accuracy.
(Nanopoulos et al, 2001)	Association Rule discovery to construct a prediction system for effective prediction of web user request. Implementation of a prefetching system.	Clarknet trace log: http://ita.ee.lbl.gov/html/traces.html	The proposed algorithm WM ₀ outperforms the existing algorithms and is very efficient in reducing the number of candidates.

the set of possible tests is limited to splitting the examples according to the value of a certain attribute. Once a node is split, the same process is performed on the new nodes, each of which contains a subset of the data in the parent node. This process is repeated until only nodes where no splits should be made remain.

Incremental induction is a method for the task of concept learning. When a new training example is entered it is classified by the decision tree. If it is incorrectly classified then the tree is revised. Restructuring the tree can be done by storing all training examples or by maintaining statistics associated with nodes in the tree.

Tree-building algorithms usually have several stopping rules. These rules are usually based on several factors including maximum tree depth, minimum number of elements in a node considered for splitting, or the minimum number of elements that must be in a new node.

The second phase of the algorithm optimizes the resulting tree obtained in the first phase. Pruning is a technique used to make a tree more general. It removes splits and the subtrees created by them. There is a great variety of different decision tree algorithms in the literature. Some of the more common algorithms are: Classification and Regression Trees (CART) (Breiman et al, 1984) (Efron et al., 1991), Chi-squared Automatic Interaction Detection (CHAID) (Kass, 1980), C4.5 (Quinlan, 1993), J4.8 (Witten and Frank, 1999), C5.0 (which implements boosting) and ID3 (Quinlan, 1986).

Rules are, at its simplest form, an equivalent form of expressing a classification tree. In order to obtain the set of rules of a decision tree, each path is traced from root node to leaf node, recording the test outcomes as antecedents and the leaf-node classification as the consequent. The process of converting a decision tree into decision rules can be done after or before pruning.

Converting a decision tree to rules before pruning has some advantages: (1) it allows distinguishing among the different contexts in which a decision node is used, (2) the pruning decision regarding an attribute test can be made differently for

each path, and (3) it improves readability. Nevertheless, obtaining the rules before pruning can produce a high number of rules. This can be solved by applying some rule-reduction techniques. The basic techniques are: (1) eliminate redundant and unnecessary antecedents, (2) use OR to combine rules in order to reduce the total number of rules to the number of dependent variables and (3) eliminate unnecessary rules. Rules can be also obtained after the tree has been pruned, producing in this case a smaller rule knowledge-base.

Algorithms such as CART, C4.5 and C5.0 include methods to generalize rules associated with a tree which removes redundancies.

2) Applications for User Modeling

In the context of user modeling, decision trees can be used to classify users and/or documents in order to use this information for personalization purposes. Decision trees can also handle noisy data and/or data with missing parameters, which makes them very useful for creating user models due to the noisy and imprecise nature of the data available. Table IV summarizes some studies and applications of Decision Trees for UM.

Classification trees are typically used to implement classification tasks. In this case the classification trees are used to construct user models used to personalize the user experience (for example regarding his/her level of expertise, his/her cognitive style, etc.). Some examples of this approach are (Tsukada et al, 2002) and (Beck et al., 2003). Due to its ability to group users with similar characteristics, classification trees can be also used to implement recommendation tasks. Examples of these approaches are (Paliouras et al, 1999), (Zhu et al., 2003) and (Webb et al., 1997).

Classification rules are widely used to model user behavior because it provides a straightforward framework to represent knowledge. The readability of the output knowledge is a great advantage of this approach. Table V summarizes some studies and applications of Classification Rules for UM.

TABLE IV
EXAMPLES OF SOME DECISION TREE-BASED USER MODELS

	Application	Input Data	Results
(Paliouras et al., 1999)	Construction of user stereotypes using C4.5. Stereotypes are used to anticipate interests of a user within the context of a news retrieval system.	Questions answered by 31 users regarding the ECRAN information system.	It is very important to have good data in order to obtain good models.
(Tsukada et al., 2001)	Automatic classification of web pages in a pre-specified set of categories using C4.5 and association rules.	14 top categories of Yahoo! JAPAN. From each category 200 pages were randomly download.	The experimental evaluation demonstrates that this method provides acceptable accuracy with the classification of web-page into top categories of Yahoo! JAPAN.
(Webb et al., 1997)	Use of C4.5 to build the Feature Based Modeling instruction module. The results are applied to the Subtraction Modeler.	Test administered to 73 nine to ten year old primary school students. Each test contained 40 three column subtraction problem randomly generated.	C4.5 increases the number of predictions made without significantly altering the accuracy of those predictions.
(Zhu et al., 2003)	Construction of a recommender system to help users find relevant information on the web using C4.5 and Naïve Bayesian Classifier.	Collected data from 129 participants, asking each participant to perform two search tasks.	C4.5 outperforms Naïve Bayesian Classifier.
(Beck et al., 2003)	Construction of a User Model for an adaptive tutor with J4.8 and Naïve Bayesian classifier.	Data collected from the interaction of 88 students with the Reading Tutor.	Naïve Bayesian Classifier outperforms J4.8 for individual modeling and J4.8 outperforms Naïve Bayesian Classifier for Group modeling.

(Shah et al., 2002), (Zhang, 2003), (Joerding, 1999) and (Semeraro et al., 2001) are examples of classification tasks implemented using classification rules. In these cases, the rules are used to characterize different user stereotypes that are used to implement some personalized feature. (Cercione, 2002) and (Cercione et al., 2002b) are examples of application of classification rules for recommendation.

Classification rules can not only be obtained from a learning algorithm but can also be directly expressed by the designer of the system. Some examples of this approach are (Morales et al, 1999) that constructs a model capturing the behavior of a user that is controlling a pole, (De Bra et al.2003) and (Romero et al., 2003) that use rules as part of the architecture of a general purpose tool for adaptive websites, and (Benaki et al., 1997) that uses rules to characterize user behavior in the context of information retrieval.

3) Limitations

Decision trees/Classification rules produce results that are highly dependent on the quality of the data available. This is caused by the fact that subtrees are created using the maximum information gain possible. In some cases if the information available is not appropriate, which typically happens when the information used to create user models has been obtained using user feedback or in a noisy environment, the models created will not correctly capture user behavior. Also, decision trees have the problem that for high dimensional problems, the response time can be very high. This is an inconvenient when working with adaptive systems, because real-time response is needed. This problem can be solved in some cases using classification rules.

Special interest for user modeling has the combination of classification rules with soft computing techniques (fuzzy logic and neural networks especially) in order to create more flexible user models. Fuzzy classification rules are able to overlap user models and to improve the interpretability of the results. So far user modeling with fuzzy classification rules has not been used at its full capacity.

B. K-Nearest Neighbor (k-NN) for User Modeling

k-Nearest Neighbor is a predictive technique suitable for classification models (Friedman et al., 1975). Unlike other learning techniques in which the training data is processed to create the model, in k-NN the training data represents the model.

1) Basic Algorithms

The simplest version is the Nearest Neighbor algorithm. In this case the training portion of the algorithm simply stores the data points of the training data $S, S=\{x_1, \dots, x_n\}$. To classify a new item i , the nearest-neighbor classifier finds the closest (according to some distance metric) training-point x_j to i and assigns it to the same class of x_j . The nearest-neighbor algorithm is traditionally implemented by computing the distance from i to every x_j of S . It can also be implemented more efficiently using Voronoi diagrams.

The k-nearest-neighbor (k-NN) algorithm is a variation of nearest-neighbor. In this case, instead of looking at only the point x_j that is closest to i , the algorithm looks at the k points in S that are closest to i . Since the class of each of the k nearest points are known, the majority category in this subset is the result of the classification. Nearest-neighbor is a special case of k-nearest-neighbor, where $k = 1$.

k-NN has two main parameters: (1) the parameter k (number of nearest cases to be used) and (2) the distance metric, d .

The selection of the metric is very important, because different metrics, used on the same training data, can result in completely different predictions. In general, the selection of the metric is application dependent. A typical metric d is the Euclidean distance, where:

$$d(\mathbf{x}, \mathbf{i}) = \sqrt{(\mathbf{x} - \mathbf{i})^T (\mathbf{x} - \mathbf{i})} \quad (18)$$

The value of k should be determined by experimentation. In general, higher values of k produce noise-proof systems, but also systems with higher computational costs. Typically k is in

TABLE V
EXAMPLES OF SOME CLASSIFICATION RULES-BASED USER MODELS

	Application	Input Data	Results
(Shah et al., 2002)	Classification rules model bid strategies in an on-line auction website	Data collected from 2 auctions over different period of time, making a total of 11,537 bids collected.	The bidder model obtained can be used to simulate online markets and/or to detect fraud in on-line auction websites.
(Zhang, 2003)	Classification of the users of an information retrieval (IR) system according to different user stereotypes.	Ratings given by 64 people with four different backgrounds (stereotypes) regarding their IR knowledge.	The classification can be effectively used to personalize the information retrieval environment.
(Joerding, 1999)	Dynamic user modeling of the interests of a web shopper. The user model represents the present interest of that particular user.	Set of visits of a user to an adaptive shopping site.	Classification Rules efficiently model dynamic human behavior.
(Cercione, 2002)	ELEM2, a novel rule induction process.	Not Described	Applications to extract user profiles that are used to construct recommender systems.
(Semeraro et al., 2001)	Use of classification rules to model the user interaction in a digital library service architecture. Comparison between the results given by C4.5 and J4.8.	Log file containing the interactions of each user with the digital library system. The data is the preprocessed and labeled.	Accurate interaction models can be inferred automatically by using J4.8. J4.8 slightly outperforms C4.5.

TABLE VI
EXAMPLES OF SOME K-NN BASED USER MODELS

	Application	Input Data	Results
(Billsus et al., 1999)	Construction of an intelligent agent designed to compile a daily news program for individual users using a user model that captures user interests using a multi-strategy approach.	Training data of the agent by 10 users during 4-8 days. K-NN is used to model Short Term interest of a user.	A multi-strategy learning approach (a short-term model and a long-term model) captures changes in user interest very efficiently.
(Schwab et al., 1999)	ELFI (Electronic Funding Information) a WWW system that provides personalized information about research funding.	ELFI log file of several months of system usage. The selection of a program is regarded as interesting for that user.	The optimum solution is obtained with a combination of K-NN and Naive Bayesian Classifiers.
(Tsiriga et al., 2002)	Initialization of the Student Model in a Algebra Tutor using k-NN and stereotypes based on recognized similarities with other students that have already interacted with the system.	Information is captured by the Algebra Tutor using a set of questions to the new user of the system. This information is used to obtain the stereotype.	The combination of K-NN and stereotypes produces very good results for Student Modeling.
(Shih et al., 2001)	Development of an on-line learning system using k-NN for providing adaptive learning materials on a course of introduction to computer networks.	Not Described	Application of k-NN compared with previous researches can retrieve more adequate materials based on previous similar cases.

the range of five to twenty rather than in the order of hundreds.

Because the model created by k-NN is basically the set of training points, this means, especially for high dimensional problems and/or for problems with a lot of training data, that the amount of memory needed to store the model is very high, and that it increases when more data is added to the system. Agglomerative Nearest Neighbor (A-NN) algorithms have been proposed to solve this problem. A-NN is designed to capture the same information as k-NN but without the necessity of storing the complete set of training data. The key idea behind A-NN is that it clusters training points that have the same class. Each class is then represented by a class-representative, against which the concept of distance is measured when classifying a new instance. A-NN has been traditionally used for collaborative clustering.

2) Applications for User Modelling

k-NN algorithms can be used very effectively for recommendation. In this case the algorithm is used to recommend new elements based on user's past behavior. A typical application is the recommendation of interesting documents. As it can be seen the concept of distance (in this case the distance between two texts) is the key element to an efficient recommendation system. Examples of k-NN used to implement recommendation tasks are (Billsus et al., 1999), (Schwab et al., 1999) and (Shih et al., 2001).

Nevertheless, like any classification algorithm, it can also be used to implement classification tasks, for example to initialize properly a new user when first visiting a system given a set of existing user models. (Tsiriga et al. 2002) is an example of k-NN used to implement a classification task. Table VI summarizes some studies and applications of k-NN for UM.

3) Limitations

k-NN is a very simple technique that produces good and straightforward results. Nevertheless the main drawbacks of this approach is the high dependence of the results from the value of k and the metric d chosen. In general, to choose both

parameters, some experimentation with the system and some knowledge on how the system works will be needed. Also, the value k should be big enough to produce a classification system noise-proof, otherwise a k-NN with a small k value will be extremely susceptible to noise. In the context of user modeling this is very important because usually the data available is noisy and imprecise due to the nature of the problem. Nevertheless, an increase in the value of k implies an increase in the computational time needed. The definition of the distance d , when used for user modeling, faces the same problems as clustering techniques: (1) the necessity of transforming user related information into a vectorial representation without losing semantic, and (2) the definition of a concept of distance among those vectors that capture the characteristics of the problem (user interest, user behavior, etc.). Another limitation of k-NN for user modeling is the response time, which is directly affected by d and by the dimensionality and number of examples of the training data. Adaptive hypermedia systems need real-time response, and k-NN may in some cases not be suitable.

C. Neural Networks for User Modeling

An Artificial Neural Network (ANN) is an information processing paradigm that is inspired by the way biological nervous systems process information. Two good introductions to Neural Networks can be found in (Fausett, 1994) and (Haykin, 1999). Although typical ANNs are designed to solve supervised problems, there are also architectures to solve unsupervised problems.

1) Basic Concepts

The key element of this paradigm is the structure of the information processing system. It is composed of a large number of highly interconnected processing elements (neurons) working in parallel. They consist of the following elements: (1) Neurones, (2) Weighted interconnections, (3) An activation rule, to propagate signals through the network and (4) learning algorithm, specifying how weights are adjusted.




Fig. 3. Architecture of an Artificial Neuron

The basic element of any ANN is an artificial neuron (Fig. 3). A neuron has N weighted input lines and a single output. The neuron will combine these weighted inputs by forming their sum and, with reference to a threshold value and activation function, it will determine its output.

Being x_1, x_2, \dots, x_N the input signals, w_1, \dots, w_N the synaptic weights, u the activation potential, θ the threshold and y the output signal and f the activation function:

$$u = \sum_{i=1}^N w_i x_i \quad (19)$$

$$y = f(u - \theta) \quad (20)$$

Defining $w_0 = \theta$ and $x_0 = -1$, the output of the system can be reformulated as:

$$y = f\left(\sum_{i=0}^N w_i x_i\right) \quad (21)$$

The activation function f defines the output of the neuron in terms of the activity level at its input. The most common form of activation function used is the sigmoid function.

There are very different ways in which a set of neurons can be connected among them. The traditional cluster of artificial neurons is called neural network. Neural networks are basically divided in three layers: The Input Layer, The Hidden Layer, which may contain one or more layers, and the output layer.

The layer of input neurons receives the data either from input files or directly from electronic sensors in real-time applications. The output layer sends information directly to

the outside world, to a secondary computer process, or to other devices such as a mechanical control system. Between these two layers can be many hidden layers. These internal layers contain many of the neurons in various interconnected structures. The inputs and outputs of each of these hidden neurons simply go to other neurons. In most networks each neuron in a hidden layer receives the signals from all of the neurons in a layer above it, typically an input layer. After a neuron performs its function it passes its output to all of the neurons in the layer below it, providing a feed forward path to the output. Another type of connection is feedback. This is where the output of one layer routes back to a previous layer.

Multi-Layer Perceptrons are the typical architecture of NNs. MLP are full-connected feed-forward nets with one or more layers of nodes between the input and the output nodes.

Classification and recognition capabilities of NNs stem from the non-linearities used within the nodes. A single-layered perceptron implements a single hyperplane. A two-layer perceptron implements arbitrary convex regions consisting of intersection of hyperplanes. A three-layer NN implements decision surfaces of arbitrary complexity (Lipman, 1987)(Looney, 1987). That is the reason why a three layer NN is the most typical architecture.

NNs learn through an iterative process of adjustments. There are two training approaches: supervised and unsupervised.

In supervised training, both the inputs and the outputs are provided. The net is trained by initially selecting small random weights and internal thresholds, and presenting all training data repeatedly. Weights are adjusted after every trial using information specifying the correct class until weights converge and the cost function is reduced to an acceptable value. The vast bulk of networks utilize supervised training. The most common supervised technique is the back-propagation learning algorithm. It uses a gradient search technique to minimize a cost function defined by the mean square error (MSE) between the desired and the actual net outputs, with l the number of training points:

$$MSE = \frac{1}{l} \sum_{i=1}^l (y_i - \hat{y}_i)^2 \quad (22)$$

The generally good performance found for the back-propagation algorithm is somewhat surprising considering that

TABLE VII
EXAMPLES OF SOME NNs BASED USER MODELS

	Application	Input Data	Results
(Bidel et al., 2003)	Classification and tracking of user navigation.	Data generated from an on-line encyclopedia.	A labeled approach to the problem produces better accuracy.
(Sas et al., 2003)	Prediction of user's next step in a virtual environment	30 users performed exploration and searching within the environment.	Very accurate predictions of the next step
(Shepperd, 2002)	Adaptive filtering system for electronic news using stereotypes.	The Halifax Herald Ltd.	Very useful for readers with specific information needs.
(Beck and Woolf., 1998)	Construction of a student model for an intelligent tutoring system.	Data collected by the tutoring system	NN-based recommendation to each individual.

it is a gradient descent technique that may find a local minimum in the cost function instead of the desired global minimum.

2) Applications for User Modeling

NNs are able to derive meaning from complicated and/or imprecise data. Also, NN do not require the definition of any metric (unlike k-NN or clustering) which make them completely application independent. No initial knowledge about the problem that is going to be solved is needed. These characteristics make NNs a powerful method to model human behaviour and an ideal technique to create user models for hypermedia applications.

NNs have been used for classification and recommendation in order to group together users with the same characteristics and create profiles and stereotypes. (Bidel et al., 2003) is an example of NNs used for classification. (Sas et al., 2003), (Beck et al., 2003), (Sheperd et al., 2002) and (Beck and Woolf, 1998) use NNs for recommendation tasks. Table VII presents more details of these applications.

3) Limitations

NNs have been successfully used for UM mainly because they do not need any heuristic to produce a model. Nevertheless it still faces important limitations. The main ones are the training time needed to produce a model (which in cases can be measured in the order of many hours and even days) and the amount of information needed. The training time is an inconvenience for creating dynamic models. Although there are techniques able to retrain NNs dynamically, the techniques used so far for UM retrain the system from scratch in case more information, e.g. a new user or a new document, is added. More research in the field of incremental learning is needed. Another important limitation of NNs is their black box behavior. While the previous techniques, to a different extent, can be interpreted and manually changed, NNs can not be interpreted which limits its applications in case a human understandable user model is needed.

D. Support Vector Machines (SVM) for User Modeling

SVM is a classifier derived from Statistical Learning Theory (Vapnick, 1995). The main advantages of SVM when used for classification problems are (1) ability to work with high dimensional data and (2) high generalization performance without the need to add a-priori knowledge, even when the dimension of the input space is very high. Excellent introductions to SVM can be found in (Cristianini et al., 2000).

The problem that SVM try to solve is to find an optimal hyperplane that correctly classifies data points and separates the points of two classes as much as possible. Figure 4 is an example of the previous situation.

1) Basic Algorithms

Given two classes, the objective is to identify the




Fig. 4. SVM maximization of the distance between classes.

hyperplane that maximizes m ,

$$m = \frac{2}{\|\mathbf{w}\|} \quad (23)$$

while at the same time classifying correctly all the examples given. Being \mathbf{w}^T the hyperplane that verifies the previous condition, all points that were part of that class will verify $\mathbf{w}^T \mathbf{x} + b > 0$, where \mathbf{x} is the point that is being validated. If a point is not part of that class, then $\mathbf{w}^T \mathbf{x} + b < 0$.

Formally the problem can be presented as follows. Given m the dimension of the problem and N the number of training points, let:

$$(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_N, y_N) \in \mathcal{R}^m \times Y, y_i \in Y, Y = \{-1, 1\} \quad (24)$$

be the set of labelled inputs, where -1 indicates that the input is not of that class and 1 indicates that the input is of that class. The decision boundary should verify:

$$\begin{aligned} \mathbf{x}_i^T \mathbf{w}^T + b &\geq 1, \forall y_i = 1 \\ \mathbf{x}_i^T \mathbf{w}^T + b &\leq -1, \forall y_i = -1 \end{aligned} \quad (25)$$

The problem can be presented as:

$$\text{Maximize } m = \frac{2}{\|\mathbf{w}\|} \quad (26)$$

$$\text{Subject to } y_i(\mathbf{x}_i \mathbf{w}^T + b) \geq 1, \forall i = 1, \dots, n$$

This formulation can be expressed as a standard quadratic problem (QP) in the form of:

$$\text{maximize } \sum_{i=1}^N \alpha_i - \frac{1}{2} \sum_{i=1, j=1}^N \alpha_i \alpha_j y_i y_j \mathbf{x}_i^T \mathbf{x}_j \quad (27)$$

$$\text{Subject to } \alpha_i \geq 0, \sum_{i=1}^N \alpha_i y_i = 0$$

In this context a global maximum α_i can always be found and \mathbf{w} can be recovered as:

$$\mathbf{w} = \sum_{i=1}^N \alpha_i y_i \mathbf{x}_i \quad (28)$$




Fig. 5. Transformation of a non-linearly separable problem into a linearly separable problem.

Many of the α_i are zero, which implies that \mathbf{w} is a linear combination of a small number of data. The set of elements \mathbf{x}_i with non-zero α_i are called support vectors. Graphically are the set of points that mark the border of the class.

This approach is valid whenever the set of points of the two classes are linearly separable. Nevertheless in real data this is usually not the case. In order to work with non-linear decision boundaries the key idea is to transform \mathbf{x}_i to a higher dimension space, so that in this new space the samples can be linearly divided (Fig. 5). The problem of this solution can be the high computational requirements. SVM solve these problems using Kernels. The relationship between the kernel function K and Φ is:

$$K(\mathbf{x}_i, \mathbf{x}_j) = \Phi(\mathbf{x}_i) \cdot \Phi(\mathbf{x}_j) \quad (29)$$

Intuitively, $K(\mathbf{x}, \mathbf{y})$ represents our desired notion of similarity between data \mathbf{x} and \mathbf{y} based on our prior knowledge of the problem. $K(\mathbf{x}, \mathbf{y})$ needs to satisfy a technical condition (Mercer condition) in order for Φ to exist. In practice, K can be directly specified, thereby specifying Φ indirectly, instead of choosing Φ . An example of a kernel function is the Gaussian kernel, which is defined as:

$$K(\mathbf{x}_i, \mathbf{x}_j) = e^{-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{std^2}} \quad (30)$$

When working with a Gaussian kernel, *std* represents the standard distribution, and ideally should represent the minimum distance between any two elements of two different classes. As it can be seen when constructing a SVM based on a Gaussian kernel, the only value that needs to be defined is

std. There are other standard kernels like polynomial, linear, or sigmoidal, each one with its own parameters. These kernels are generic in the sense that they can, in some way or another, be applied to all problems and will provide good solutions. For example Gaussian kernels tend to work very well in image recognition problems. Nevertheless, results can always be improved by designing a problem-specific kernel, although this approach requires an a priori knowledge about the problem.

When working with kernels, in general it would not be possible to obtain \mathbf{w} . Nevertheless SVM can be still be used. Being N_s the number of support vectors of the training set, the decision function can be expressed as:

$$f(x) = \text{Sign} \left(\sum_{i=1}^{N_s} \alpha_i y_i K(\mathbf{x}_i, \mathbf{x}) + b \right) \quad (31)$$

Although the theoretical background given has introduced only classification system for only two classes, SVM can be generalized to a set of C classes. In this case each one of the classes will be trained against the rest $C-1$ classes, reducing the problem to a 2-class classification problem.

2) Applications for User Modelling

Due to its ability to work extremely well with high dimensional data, SVM can be used to classify documents in order to personalize recommendations given by search engines. (Liao, 2002), (Ruvini, 2003) and (Aihara et al., 2001) are examples of classification tasks implemented using SVM. Table VIII summarizes some studies and applications of SVM for UM.

3) Limitations

In general, the main problem that SVM faces is the selection of the kernel K used to implement the classifier. Although there are a set of standard kernels (polynomial, Gaussian, etc.) that work well for many applications, for complex cases it is convenient to develop a specific kernel considering the characteristics of the problem. For example, there are kernels specifically designed for gene-detection or for text classification.

In the context of user modelling, SVM has been used to personalize recommendation applications. So far the

TABLE VIII
EXAMPLES OF SOME SVM-BASED USER MODELS

	Application	Input Data	Results
(Liao, 2002)	Profiling user behavior in Windows NT using CPU time to detect a misuse or an attack on the system.	Goldring's Windows NT Profiling dataset. A user login is then transformed into a vector containing CPU time of each process.	The system identifies correctly each user if the set of users is small (5 users). For bigger sets of users (10 users) the classification rate is not as satisfactory.
(Ruvini, 2003)	A front-end to Google search engine that uses SVM to infer the user's goal and filter the set of links in order to avoid the problem of screen size in mobile devices.	Recorded users interactions with Google in a predetermined set of searching tasks.	The problem of small screen sizes is satisfactorily solved.
(Aihara et al., 2001)	Interactive document retrieval system supported by a personalization of a basic classification of categories that is adapted to each user using SVM.	Not Described	SVM effectively adapt the boundaries of each category according to each user interests.

TABLE IX
SELECTION OF SUITABLE DATA MINING TECHNIQUE

	Labeled Data		Unlabelled Data	
Task	Readability Needed	Readability Not Needed	Readability Needed	Readability Not Needed
Recommendation	Decision Trees Classification Rules	Neural Networks k-NN	Association Rules SOM	K-means Clustering Fuzzy Clustering
Classification	Decision Trees Classification Rules	SVM Neural Networks k-NN	SOM	K-means Clustering Fuzzy Clustering

classification has been applied to the documents but not to the users directly. Due to the ability of SVM to optimize the separation between the classes, the application of SVM to create classes of users can improve the individual degree of satisfaction of each user. More research is needed in the applications of SVM for UM, ranging from general applications to the development of specific kernels.

V. CRITERIA FOR THE SELECTION OF THE TECHNIQUES

The preceding discussion has demonstrated that each technique captures different relationships among the data available and expresses it using different data structures. In this section we present guidelines to help decide which technique to use when developing an adaptive hypermedia application.

We consider that, in the context of UM, there are three main criteria that determine which data mining/machine learning technique is suitable for a specific adaptive application: (1) the labeled/unlabelled nature of the data

available; (2) the type of task that is going to be implemented (Recommendation or Classification) and (3) the “readability” needed for the results. “Readability” is defined as the ability of a technique to produce a human-readable output of the knowledge captured for a non-technical user. There are two possible values for Readability: (1) needed and (2) not needed. The first one expresses the necessity of having a human readable output while the second one states that this factor is not relevant. Table IX presents a guideline of what data mining techniques are useful considering the criteria previously introduced. The techniques are classified according to the set of references used in this study. The set of techniques that can be applied when the systems needs readability can also be applied when this factor is not relevant.

In general, when selecting a data mining technique two of the more important factors are (1) the ability to handle high dimensional data and (2) scalability. Although the ability to handle high dimensional data for a generic problem is a very important characteristic, within the context of UM we consider that it is not a very relevant factor because of the

TABLE X
GENERAL CHARACTERISTICS OF THE REVISED TECHNIQUES

	Off-Line Complexity (Indication of Scalability)	Dynamic Modeling	Labeled / Unlabeled	Readability
K-means Clustering	$O(k \cdot m \cdot n \cdot i)$ (Hartigan, 1975) n number of instances to cluster m number of attributes k number of clusters i number of iterations, with $i=O(n)$ (Davidson et al., 2003)	No	Unlabeled	No
SOM	$O(n)$, with n the number of feature vectors (Ramsey et al., 1999) $O(M^2)$ where M is the number of map units: each learning step requires $O(M)$ and to achieve a sufficient statistical accuracy the number of iterations should be at least some multiple of M . (Kaski, 1997)	No	Unlabeled	Yes
Fuzzy Clustering	$O(n^2)$ with n the number of objects For some optimized algorithms $O(n \log n)$ (Krishnapuram et al., 2001)	No	Unlabeled	No
Association Rules	NP-Complete Exponential with the number of items (Angiulli et al., 1998)	No	N/A	Yes
Decision Trees	For single attribute, multi-way splits on A discrete variables and data size of N : $O(A^2 N)$ For continuous attributes: $O(A^2 N^2)$ (Martin et al., 1996) Pruning: $O(N \cdot h)$ (Martin et al., 1996)	Yes	Labeled	Yes
Classification Rules	Same as Decision Trees + Rule generation	Yes	Labeled	Yes
k-NN	Linear with the number of samples Empirical sample complexity is exponential in the number of irrelevant features (Langley et al. 1993)	Yes	Labeled	No
Neural Networks	NP-Complete for a generic 3 layer NN (Blum et al., 1992) Polynomial for some simple two layer networks (Blum et al., 1992)	Yes	Both	No
SVM	Complexity of Solving a Quadratic Optimization problem (QP) at each iteration: $O(N^3)$ with N total number of training points. In general it is highly dependent of the SVM implementation used.	No	Labeled	No

dimensionality of the data used. Nevertheless, in the context of user modeling, the scalability of the techniques is a very important factor due to the high number of users that, in general, will interact with an adaptive hypermedia system. The scalability of each technique regarding the number of users will depend on how the information of each user is presented. An indication of the scalability of each technique is presented in the first column of Table X.

Table X summarizes the characteristics of the techniques presented along four dimensions. The first three dimensions capture some of the main problems that machine learning for user modeling faces according to (Webb et al., 2001): Computational Complexity for off-line processing; Dynamic Modeling, which indicates the suitability of the technique to change a user model on-the-fly; and Labeled/Unlabeled. The "Readability" dimension has also been added to the table.

The combination of Table IX and Table X can be used to guide a choice of which technique to use when modeling user behavior for adaptive systems. First, Table IX identifies the set of techniques suitable for the adaptive application and, after that, Table X can be used to refine the choice considering the scalability and dynamic modeling capabilities of each technique.

Finally, there are other techniques not reviewed in this paper, mainly predictive statistical techniques (Zukerman et al., 2001), that can be also used to create user models. For example, recommendation tasks have also been implemented with Markov models, like in (Albrecht et al., 1999), (Deshpande et al., 2001), (Duchamp, 1999), (Sarukkai, 2000) (Anderson et al., 2001) and (Anderson et al., 2002) or with Bayesian networks (Jensen, 1996), like in (Conati et al., 1997), (Grobmann-Hutter et al., 1997), (Witting, 2003), (Horvitz et al., 1998) and (Schaffer et al., 1997). Classification tasks have also been implemented with Bayesian Classifiers, like in (Chen et al., 2002) and (Schwab et al., 1999).

VI. CONCLUSION AND DISCUSSION

Hypermedia systems are becoming more important in our everyday activities and their contents and services are evolving every day. Due to this important role, users welcome any improvement on the hypermedia services they receive. One trend used to improve digital services is through personalization which is based in the concept of user modeling. In this context data mining and machine learning techniques can be used to automatically identify user patterns and interests and use the knowledge obtained to produce user models.

This paper has presented a review of the state of the art of data mining techniques within the area of user modeling for adaptive hypermedia systems. The review demonstrates that one of the main problems that the development of user model faces is the lack of any kind of standardization for the design of user models. In order to improve this situation this paper has tried to give a set of guidelines that formalize the design of user models using a data mining approach. In the same line

of standardization, another interesting research area would be the design of a system to evaluate the performance of user modeling in order to make it feasible to compare different user models.

It is our opinion that the future of UM is in hybrid systems. As has been shown, each technique captures different elements of user behavior. The combination of these techniques among themselves and with other machine learning techniques, especially with soft computing techniques, will provide a useful framework to efficiently capture the natural complexity of human behavior.

Acknowledgments

The work presented in this paper is funded by the UK Arts and Humanities Research Board (AHRB grant reference: MRG/AN9183/APN16300).

REFERENCES

- [1] Agrawal, R., Imielinski, T., Swami, A., Mining Association Rules between Sets of Items in Large Databases in Proceedings of the 1993 ACM SIGMOD Conference, Washington DC, USA, 1993
- [2] Aihara, K., Takasu, A.: Category Based Customization Approach for Information Retrieval, Proceedings of the 8th International Conference on User Modeling, LNAI 2109, 2001
- [3] Albrecht, D., Zukerman, I., Nicholson, A., Pre-sending documents on the WWW: A comparative study. IJCAI99-Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, 1999
- [4] Anderson, C. R., Domingos, P., Weld, D.S., Adaptive Web Navigation for Wireless Devices, in Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-01), 2001
- [5] Anderson, C., Domingos, P., Weld, D.: Relational Markov Models and their Application to Adaptive Web Navigation". In Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2002), 143-152 (2002)
- [6] Angiulli, F., Ianni, G., and Palopoli, L., "On the complexity of mining association rules", Data Mininf and Knowledge Discovery (2):3, pp. 263-281 (1998)
- [7] Beck, J., Jia, P., Sison, J., Mostow, J.: Predicting Student Help-Request Behavior in an Intelligent Tutor for Reading, Proc. of the 9th Int. Conf. on User Modeling, LNAI 2702, pp. 303-312, 2003
- [8] Beck, J.E., Woolf, B.P.: Using a Learning Agent with a Student Model, Lecture Notes in Computer Science, LNCS 1452, 1998
- [9] Benaki, E., A. Karkaletsis, V.A., Spyropoulos, C.D.: User Modeling in WWW: the UMIE Prototype, Adaptive Systems and User Modeling on the World Wide Web" Sixth International Conference on User Modeling, Chia Laguna, Sardinia, 2-5 June 1997
- [10] Bezdek, J.C.: Pattern Recognition with Fuzzy Objective Function Algorithms, Plenum Press, New York, 1981
- [11] Bidel, S., Lemoine, L., Piat, F., Artieres, T., Gallinari, P., "Statistical machine learning for tracking hypermedia user behavior", Second Workshop on Machine Learning, Information Retrieval and User Modelling, 9th International Conference in User Modeling, Pittsburgh, PA, 2003
- [12] Bidel, S., Lemoine, L., Piat, F.: Statistical machine learning for tracking hypermedia user behavior, 2nd Workshop on Machine Learning, Information Retrieval and User Modeling, 9th Int. Conf. in UM, 2003
- [13] Billsus, D., Pazzani, N.: A hybrid User Model for News Story Classification. In Proceedings of the 7th International Conference on User Modelling, Banff, Canada, 1999, 99-108
- [14] Billsus, D., Pazzani, N.: A hybrid User Model for News Story Classification. In Proceedings of the 7th International Conference on User Modelling, Banff, Canada, 1999, 99-108 Billsus, D. Pazzani, N.: Learning Probabilistic User Models, Proceedings of the Workshop on Machine Learning for User Modeling, 1997
- [15] Blum, A.L., Rivest, R.L., "Training a 3-Node Neural Network is NP-Complete", Neural Networks (5):1, pp117-127, 1992

- [16] Boser, B. E., Guyon, I. M., Vapnik, V. N., 1992. A training algorithm for optimal margin classifiers. Proceedings of the fifth annual workshop on Computational learning theory 144-152
- [17] Breiman, L., Friedman, J.H., Olshen, R.A., and Stone, C.J.: *Classification and Regression Trees*, Wadsworth Int. Group/Probability Series, Belmont, California, USA, 1984
- [18] Brusilovsky, P., Schwarz, E.: User as Student: Towards an Adaptive Interface for Advanced Web-Based Applications. In A. Jamesson, C. Paris and C. Tasso (Eds.), *User Modeling: Proceedings of the Sixth International Conference, UM97*.
- [19] Callan, J., Smeaton, A., Beaulieu, M., Borlund, P., Brusilovsky, P., Chalmers, M., Lynch, C., Riedl, J., Smyth, B., & Straccia, U. (2001). Personalization and Recommender Systems in Digital Libraries. *Proc. of the 2nd DELOS Workshop on Personalization and Recommender Systems in Digital Libraries*.
- [20] Cercone, N., Hou, L., Keselj, V., An, A., Naruedomkul, K., Hu, X.: "From Computational Intelligence to Web Intelligence", IEEE Computer, November 2002, 72-76, 2002
- [21] Cercone, N.: From Computational Intelligence to Web Intelligence: An Ensemble from Potpourri. In *Web Intelligence: Research and Development, LNAI 2198*, 31-42, 2002
- [22] Chen A., Lin F., Ma, W., Wenyin, L.: User Intention Modeling in Web Applications Using Data Mining, *World Wide Web: Internet and Web Information Systems*, 5, 181-191, 2002, Kluwer
- [23] Conati, C., Gertner, A., Vanlehn, K., Druzdzel, M.: On-line student modelling for coached problem solving using Bayesian Networks. In *Proceedings of the 1997 User Modelling Conference, Sardinia, 1997*
- [24] Cristianini, N., Shawe-Taylor, J., 2000. *An Introduction to Support Vector Machines*. Cambridge University Press.
- [25] Davidson, I., Satyanarayana, A., "Speeding up k-means Clustering by Bootstrap Averaging", IEEE Data Mining Workshop on Clustering Large Data Sets, Third IEEE International Conference on Data Mining, 2003
- [26] DeBra, P., Stash, N.: AHA! A General-Purpose Tool for Adaptive Websites, 2002 World Wide Web Conference, May 2002
- [27] Deshpande, M., Karypis, G.: Selective Markov Models for predicting web-page accesses. In *Proceedings SIAM Int. Conference on Data Mining (SDM2001)*, 2001
- [28] Doux, A., Laurent, J., Nadal, J.: Symbolic Data Analysis with the k-Means Algorithm for User Profiling, *User Modeling: Proceedings of the sixth International Conference, UM97*, 1997
- [29] Duchamp, D., Prefetching Hyperlinks, In *Proc. Second USENIX Symp. on Internet Technologies and Systems, USENIX, Boulder, CO, October 1999*, pp. 127-138
- [30] Efron, B., Tibshirani, R.: *Statistical Data Analysis in the Computer Age*, Science, Vol. 253, No. 5018, 390-395, 1991
- [31] Everitt, B.S.: *Cluster Analysis*, Arnold, 3rd edition, 1993
- [32] Fausett L., *Fundamentals of Neural Networks*, Prentice-Hall, 1994
- [33] Fink, J., Kobsa, A., Nill, A.: Adaptable and Adaptive Information Access for All Users, Including the Disabled and the Elderly. In A. Jamesson, C. Paris and C. Tasso (Eds.), *User Modeling: Proceedings of the Sixth International Conference, UM97*.
- [34] Ford, N., Chen, S.: Individual Differences, Hypermedia Navigation and Learning: An Empirical Study. In *Journal of Education Multimedia and Hypermedia*, Vol. 9 (4), pp. 281-311, 2000
- [35] Frias-Martinez, E., Karamcheti, K., "A Prediction Model for User Access Sequences", in *Proceedings of the 4th WEBKDD Web Mining for Usage Patterns and User Profiles Workshop part of The Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 23-26 July 2002, Edmonton, Canada, pp. 53-62
- [36] Friedman, J.H., Baskett, F., Shustek, L.J.: An algorithm for finding nearest neighbors, *IEEE Transactions on Computers*, 24:1000-1006, 1975
- [37] Fu, Y., Sandhu, K., and Shih, M.-Y. (1999). Clustering of Web users based on access patterns. In *Proceedings of the 1999 KDD Workshop on Web Mining*, San Diego, CA. Springer-Verlag
- [38] Geyer-Schulz, A., Hahsler, M.: Evaluation of Recommender Algorithms for an Internet Information Broker based on Simple Association Rules and on the Repeat-Buying Theory. In *Proceedings of WEBKDD 2002, Fourth WEBKDD Web Mining for Usage Patterns&User Profiles*, Edmonton, Canada, 100-114, 2002
- [39] Goren-Bar, D., Kuflik, T., Lev, D., Shoval, P.: Automatic Personal Categorization Using Artificial Neural Networks, 8th International Conference on User Modeling 2001, LNAI 2109, pp. 188-198, 2001
- [40] Grobmann-Hutter, B., Jameson, A., Witting, F.: Learning Bayesian Networks with Hidden Variables for User Modelling. In *Proceedings of the Workshop Machine Learning for User Modelling, 6th International Conference on User Modeling, 1997*
- [41] Hartigan, J., *Clustering Algorithms*, Wiley Publishing, 1975
- [42] Hay, B, Wets, G., Vanhoof, K.: Clustering navigation patterns on a website using a Sequence Alignment Method, *IJCAI 2001 Workshop on Intelligent Techniques for Web Personalization*
- [43] Haykin S., *Neural Networks*, 2nd Edition, Prentice Hall, 1999
- [44] Horvitz, E., Breese, J., Heckerman, D., Hovel, D., Rommelse, K.: The Lumière project: bayesian user modeling for inferring the goals and needs of software users. *Uncertainty in Artificial Intelligence: Proceedings of the Fourteenth Conference* (G. Cooper, S. Moral, eds.) Morgan Kaufmann, San Francisco, 256-265, 1998
- [45] J.K. Martin and D.S. Hirschberg, "On the complexity of learning decision trees," *Proc. 4th Int. Symp. on Artif. Intell. and Math.*, Fort Lauderdale, FL (1996) 112-115.
- [46] Jain, A., Dubes, R.C.: *Algorithms for Clustering Data*, Prentice Hall, 1988
- [47] Jain, A.K., Dubes, R.C.: *Data clustering*, ACM Computing Survey, 31:264-323, 1999
- [48] Jensen, F.V.: *An Introduction to Bayesian Networks*, Springer-Verlag, 1996
- [49] Joerding, T., Temporary User Modeling Approach for Adaptive Shopping on the Web, *Proceedings of the 2nd Workshop on Adaptive Systems and User Modeling on the WWW*, 1999
- [50] Joshi, A., Joshi, K., Krishnapuram, R., On mining Web Access Logs, in *Proceedings of the ACM-SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*, 2000, pp. 63-69
- [51] Joshi, A., Joshi, K., Krishnapuram, R., On mining Web Access Logs, in *Proceedings of the ACM-SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*, 2000, pp. 63-69
- [52] Kaski, S., "Data Exploration Using Self Organizing Maps", *Acta Polytechnica Scandinavica, Mathematics, Computing and Management in Engineering Series No. 82*, 1997
- [53] Kass, G.V., *An Exploratory Technique for Investigating Large Quantities of Categorical Data*, *Applied Statistics* 29: p. 119-127, 1980
- [54] Kobsa, A.: Generic User Modeling Systems, *User Modeling and User-Adapted Interaction* 11, 49-63, 2001
- [55] Kohonen, T., *Self-Organizing Maps*, New York : Springer-Verlag, 1997
- [56] Krishnapuram, R., Joshi, A., Nasraoui, O., Yi, L., Low-Complexity Fuzzy Relational Clustering Algorithms for Web Mining, in *IEEE Transactions on Fuzzy Systems*, August 2001, Vol. 9 (4), pp. 595-608, 2001
- [57] Lampinen, T., Koivisto, H.: Profiling Network Applications with Fuzzy C-Means Clustering and Self-Organising Map, *International Conference on Fuzzy Systems and Knowledge Discovery*, November 2002
- [58] Langley, P. and Iba, W., "Average-case analysis of a nearest neighbor algorithm". *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, pp.889-894, 1993
- [59] Liao, Y.: Windows NT User Profiling with Support Vector Machines, *Proceedings of the 2002 UC Davis Student Workshop on Computing*, Technical Report CSE-2002-28, Dept. Computer Science, UC Davis 2002
- [60] Lin, F., Liu Wenyin, Zheng Chen, Hongjiang Zhang, Tang Long: User Modeling for Efficient Use of Multimedia Files, *Advances in Multimedia Information Processing - PCM 2001 : Second IEEE Pacific Rim Conference on Multimedia Beijing, China, October 24-26, 2001*
- [61] Lippmann, R.P., 1987. An Introduction to Computing with Neural Nets, *IEEE ASSP Magazine*, April 1987, pp 4-28.
- [62] Looney, C.G., *Pattern Recognition using Neural Networks: Theory and Algorithms for Engineers and Scientists*, Oxford University Press, 1997
- [63] MacQueen, J.B.: Some Methods for classification and Analysis of Multivariate Observations, *Proceedings of 5-th Berkeley Symposium on Mathematical Statistics and Probability*, Berkeley, University of California Press, 1:281-297, 1967
- [64] Magoulas, G., Papanikolaou, K., Grigoriadou, M.: Adaptive web-based learning: accomodating individual differences through system's

- adaptation. In *British Journal of Educational Technology*, Vol 34 (4), pp. 1-19, 2003
- [65] Martin, J.K. and Hirschberg, D.S., "The time complexity of Decision Tree Induction", Department of Information and Computer Science, University of California at Irvine, Technical Report 95-27 (ICS/TR-95-27), 1995
- [66] Mitchell, T. "Decision Tree Learning", *Machine Learning*, The McGraw-Hill Companies, Inc., 1997, pp. 52-78.
- [67] Mobasher, B. Jain, N., Srivastava, J.: *Web Mining: Patterns from WWW Transactions*. Technical Report TR97-050, Department of Computer Science, University of Minnesota, March
- [68] Mobasher, B. Jain, N., Srivastava, J.: *Web Mining: Patterns from WWW Transactions*. Technical Report TR97-050, Department of Computer Science, University of Minnesota, March 1997
- [69] Mobasher, B., Cooley, R., Automatic Personalization Based on Web Usage Mining, in *Communications of the ACM*, vol. 43, num. 8, 2000, pp.142-151
- [70] Mobasher, B., Dai, H., Luo, T., Nakagawa, M.: Effective personalization based on Association Rule discovery from web usage data. In 3rd. ACM Workshop on Web Information and Data Management, 2001.
- [71] Morales, R., Pain, H.: Modelling of Novices' Control Skills with Machine Learning, *User Modeling: Proceedings of the Seventh International Conference, (UM99)*. June 1999 ,159-168
- [72] Nanopoulos, A., Katsaros, D., Manolopoulos, Y., Effective Prediction of Web-user Accesses: A Data Mining Approach, *Proceeding of the WEBKDD 2001 Workshop*, San Francisco, CA, 2001
- [73] Nasraoui, O., Frigui, H., Joshi, A., Krishnapuram, R.: Mining Web Access Logs Using Relational Competitive Fuzzy Clustering, *Eight International Fuzzy Systems Association World Congress - IFSA 99*, Taipei, August 99
- [74] Paliouras, G., Karkaletsis, V., Papatheodorou, C., Spyropoulos, C.: Exploiting Learning Techniques for the Acquisition of User Stereotypes and Communities, *Proceedings of the International Conference on User Modelling (UM '99)*, 1999
- [75] Paliouras, G., Papatheodorou, C., Karkaletsis, V., Spyropoulos, C.: Clustering the users of large web sites into communities. In 17th International Conference on Machine Learning, 719-726, 2000
- [76] Perkowitz, M., Etzioni, O.: Adaptive Web Sites. *Communications of the ACM*, Vol. 43, No. 8, 152-158, 2000
- [77] Perkowitz, M., Etzioni, O.: Adaptive Web Sites: an AI Challenge. In *IJCAI*, 16-23, Nagoya, Japan, 1998
- [78] Perkowitz, M., Etzioni, O.: Towards Adaptive Web Sites: Conceptual Framework and Case Study. In *The Eighth International World Wide Web Conference*, 1999
- [79] Quinlan, J.R., *Induction of Decision Trees*. *Machine Learning* (1), p. 81-106, 1986
- [80] Quinlan, J.R.: *C4.5: Programs for Machine Learning*. Morgan Kaufmann, 1993
- [81] Quiroga, L.M., Mostafa, J.: Empirical evaluation of explicit versus implicit acquisition of user profiles. In *Proceedings of the fourth ACM conference on Digital libraries*, pp. 238-239, 1999.
- [82] Ramsey, M., Chen, H. and Zhu, B. "A Collection of Visual Thesauri for Browsing Large Collections of Geographic Images". *Journal of the American Society for Information Science & Technology* 50(9):826-834, 1999
- [83] Riecken, D. (2000). Personalized Views of Personalization. *Communications of the ACM*, Vol. 43 (8), 27-28.
- [84] Romero, C., Ventura, S., de Bra, P. de Castro, C.: Discovering Prediction Rules in AHA! Courses, *Proceedings of the 9th International Conference on User Modeling*, LNAI 2702, 2003-11-21
- [85] Romero, C., Ventura, S., de Bra, P., de Casto, C.: Discovering Prediction Rules in AHA! Courses, *Proceedings of the 9th International Conference on User Modeling*, LNAI 2702, 2003-11-21
- [86] Ruvini, J.: Adapting to the User's Internet Search Strategy, , *Proceedings of the 9th International Conference on User Modeling*, LNAI 2702, 2003
- [87] Sarukkai, R.R.: Link Prediction and Path Analysis using Markov Chains. In *Computer Networks*, 33(1-6):377-386 (2000)
- [88] Sas, C., Reilly, R., O'Hare, G.: A Connectionist Model of Spatial Knowledge Acquisition in a Virtual Environment. In 2nd Workshop on Machine Learning, Information Retrieval and User Modeling, 2003
- [89] Schafer, R., Weyrath, T.: Assessing Temporally Variable User Properties with Dynamic Bayesian Networks, *international Conference in User Modeling*, UM97, 1997
- [90] Schwab, I., Pohl, W.: Learning Information Interest from Positive Examples. In *Proceedings of the International Conference on Machine Learning & Applications*, pp. 15-20. Chania, Greece, 1999
- [91] Semeraro, G., Ferilli, S., Fanizzi, N., Abbattista, F.: Learning Interaction Models in a Digital Library Service, *Proceedings of the 8th International Conference on User Modeling*, LNAI 2109, 2001
- [92] Shah, H.S., Joshi, N.R., Wurman, P.R.: Mining for Bidding Strategies on eBay. In *Proceedings of WEBKDD 2002*, Fourth WEBKDD Web Mining for Usage Patterns & User Profiles, Edmonton, Canada, 16-30, 2002
- [93] Sheperd, A., Watters, C., Marath, A.T.: Adaptive User Modeling for Filtering Electronic News. In *Proc. of the 35th Annual Hawaii Intl. Conf. on System Sciences (HICSS-02)*, Vol. 4, 2002
- [94] Shih, B. Lee, W.: The Application of nearest neighbour Algorithm on Creating an Adaptive On-Line Learning System, in 31st ASEE/IEEE Frontiers in Education Conference, October 10-13, 2001 Reno, NV
- [95] Tsiriga, V., Virvou, M. Initializing the Student Model using Stereotypes and Machine Learning *Proceedings of the 2002 IEEE International Conference on Systems, Man and Cybernetics*, 2002
- [96] Tsukada, M., Washio, T., Motoda, H., "Automatic Web-Page Classification by Using Machine Learning Methods", in *Web Intelligence: Research and Development*, LNAI 2198, 303-313, 2001
- [97] Vapnik, V., 1995. *The nature of statistical learning theory*. Springer-Verlag, New York
- [98] Webb, A.: *Statistical Pattern Recognition*, Arnold, 1999
- [99] Webb, G. I., B. C. Chiu, and M. Kuzmycz (1997). Comparative Evaluation of Alternative Induction Engines for Feature Based Modelling. *International Journal of Artificial Intelligence in Education*, 8:97-115, 1997
- [100] Webb, G.I., Pazzani, M.J., Billsus, D.: Machine Learning for User Modeling. In *User Modeling an User-Adapted Interaction*, (11): 19-29, 2001
- [101] Wijnen, J., Meersman, R., On the Complexity of mining association rules. *Data Mining and Knowledge Discovery* (2):3, pp. 263-281, 1998
- [102] Winston, P."Learning by Building Identification Trees", *Artificial Intelligence*, Addison-Wesley Publishing Company, 1992, pp. 423-442
- [103] Witten, I.H., Frank E.: *Data Mining. Practical Machine Learning Tools and Techniques with JAVA Implementations*. Morgan Kaufman Publishers, 1999
- [104] Witting, F.: ML4UM for Bayesian Network User Model Acquisition. In *Second Workshop on Machine Learning, Information Retrieval and User Modeling*, June 22-26, Pittsburgh, PA, USA, 2003
- [105] Zhang, X.: Discriminant Analysis as a Machine Learning Method for Revision of User Stereotypes of Information Retrieval Systems, *MLIRUM'03: Second Workshop on Machine Learning, Information Retrieval and User Modeling at the Ninth International Conference on User Modeling*, June 22-26, Pittsburgh, PA, USA
- [106] Zhu, T., Greiner, R., Haubl, G.: Learning a Model of a Web User's Interests, *Proceedings of the 9th International Conference on User Modeling*, LNAI 2702, 2003-11-21
- [107] Zukerman, I., Albrecht, D.W., "Predictive Statistical Models for User Modeling", *User Modeling and User-Adapted Interaction* 11: pp.5-18, 2001.
- [108] Zukerman, I., Albrecht, D.W., Nicholson, A.E., "Predicting Users Request on the WWW", *Proceedings of the 7th International Conference on User Modeling*, UM99 Conference