

EXHIBIT 1



US006981040B1

(12) **United States Patent**
Konig et al.

(10) **Patent No.:** **US 6,981,040 B1**

(45) **Date of Patent:** **Dec. 27, 2005**

(54) **AUTOMATIC, PERSONALIZED ONLINE INFORMATION AND PRODUCT SERVICES**

(75) Inventors: **Yochai Konig**, San Francisco, CA (US);
Roy Twersky, San Francisco, CA (US);
Michael R. Berthold, Berkeley, CA (US)

(73) Assignee: **Utopy, Inc.**, San Francisco, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 914 days.

(21) Appl. No.: **09/597,975**

(22) Filed: **Jun. 20, 2000**

Related U.S. Application Data

(60) Provisional application No. 60/173,392, filed on Dec. 28, 1999.

(51) **Int. Cl.**⁷ **G06F 15/173**

(52) **U.S. Cl.** **709/224; 709/223; 709/228; 715/736**

(58) **Field of Search** 709/200, 201–203, 709/223–225, 27–228; 707/1–3, 7–10, 101; 715/500, 736, 513–514

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,704,017 A	12/1997	Heckerman et al.	395/61
5,754,939 A *	5/1998	Herz et al.	455/3,04
5,867,799 A	2/1999	Lang et al.	707/1
5,918,014 A	6/1999	Robinson	395/200.49
5,933,827 A	8/1999	Cole et al.	707/10
5,964,839 A *	10/1999	Johnson et al.	709/224
5,983,214 A	11/1999	Lang et al.	707/1
5,991,735 A *	11/1999	Gerace	705/10
5,999,975 A	12/1999	Kittaka et al.	709/224
6,006,218 A *	12/1999	Breese et al.	707/3

(Continued)

OTHER PUBLICATIONS

Mobasher, B., *Automatic personalization based on web usage mining*, <http://maya.cs.depaul.edu/~mobasher/personalization>.

Yann LeCun (2004) in a lecture entitled “Machine Learning and Pattern Recognition” and presented at The Courant Institute, New York University.

Basset et al. (2002) in a paper entitled “A Study of Generalization Techniques in Evolutionary Rule Learning”.

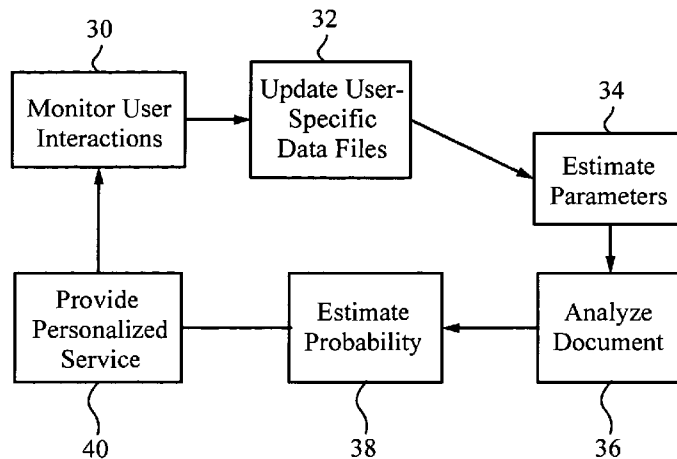
Primary Examiner—Bharat Barot

(74) *Attorney, Agent, or Firm*—Lumen Intellectual Property Services, Inc.

(57) **ABSTRACT**

A method for providing automatic, personalized information services to a computer user includes the following steps: transparently monitoring user interactions with data during normal use of the computer; updating user-specific data files including a set of user-related documents; estimating parameters of a learning machine that define a User Model specific to the user, using the user-specific data files; analyzing a document to identify its properties; estimating the probability that the user is interested in the document by applying the document properties to the parameters of the User Model; and providing personalized services based on the estimated probability. Personalized services include personalized searches that return only documents of interest to the user, personalized crawling for maintaining an index of documents of interest to the user, personalized navigation that recommends interesting documents that are hyperlinked to documents currently being viewed; and personalized news, in which a third party server customized its interaction with the user. The User Model includes continually-updated measures of user interest in words or phrases, web sites, topics, products, and product features. The measures are updated based on both positive examples, such as documents the user bookmarks, and negative examples, such as search results that the user does not follow. Users are clustered into groups of similar users by calculating the distance between User Models.

62 Claims, 19 Drawing Sheets



US 6,981,040 B1

Page 2

U.S. PATENT DOCUMENTS

6,029,161 A	2/2000	Lang et al.	707/1	6,567,850 B1 *	5/2003	Freishtat et al.	709/224
6,041,311 A	3/2000	Chislenko et al.	705/27	6,687,696 B2	2/2004	Hofmann et al.	
6,564,170 B2 *	5/2003	Halabieh	709/224	6,732,090 B2 *	5/2004	Shanahan et al.	715/500

* cited by examiner

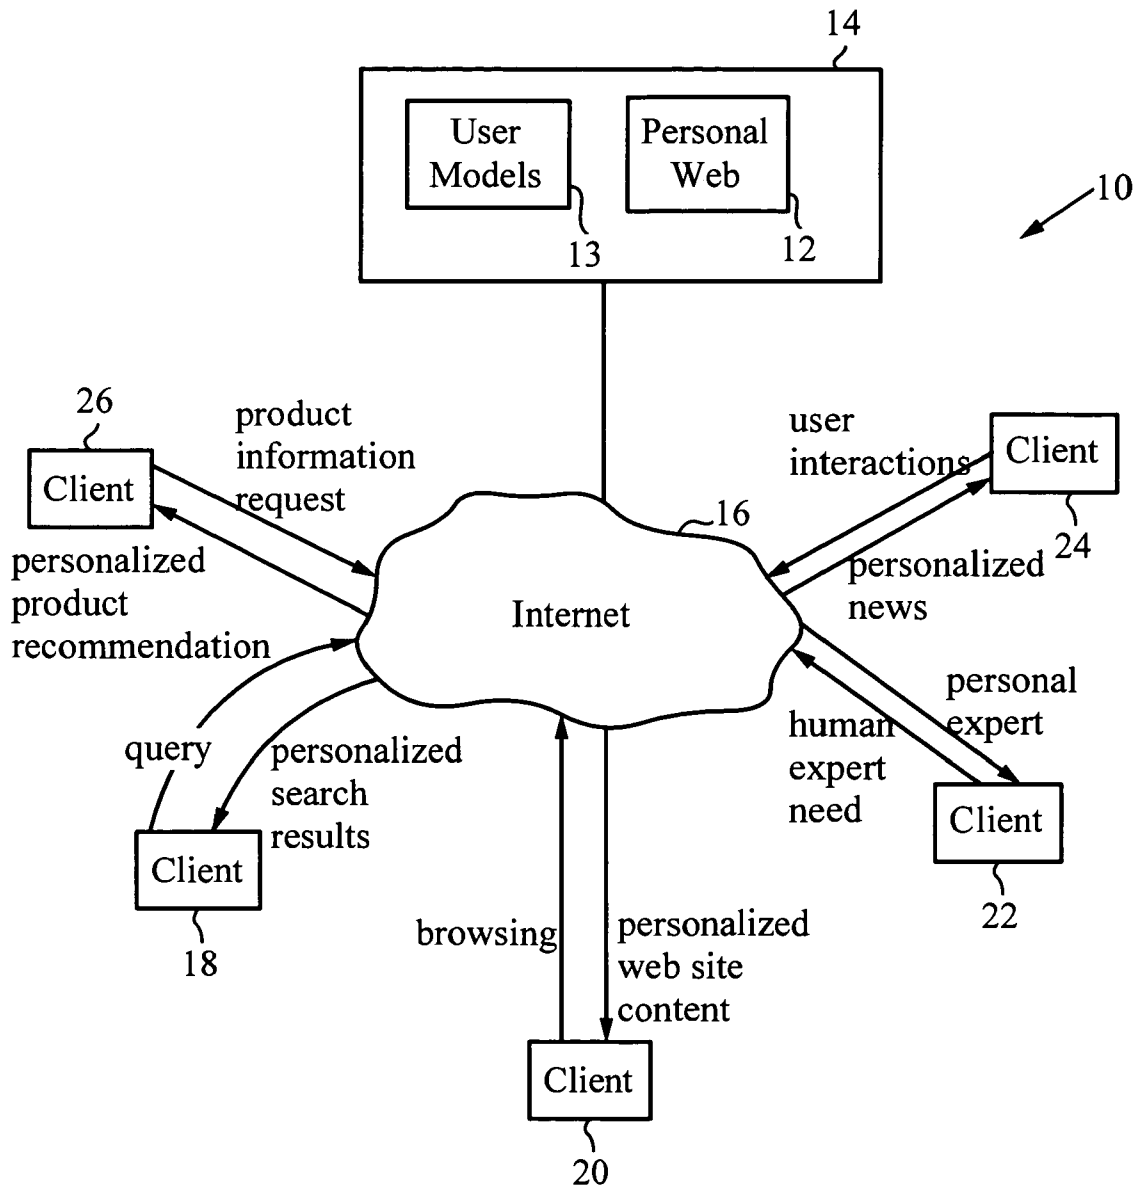


Fig. 1

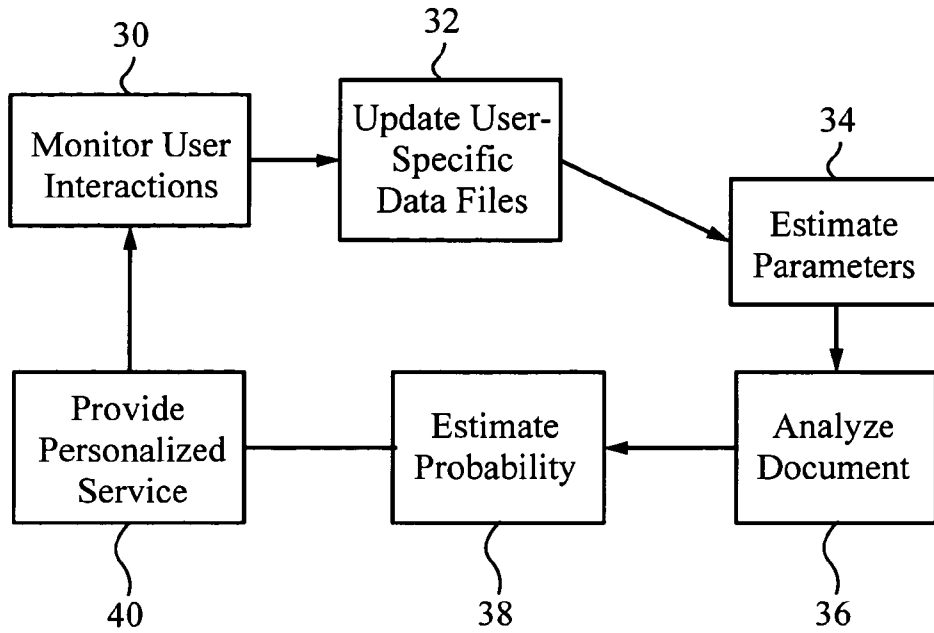


Fig. 2

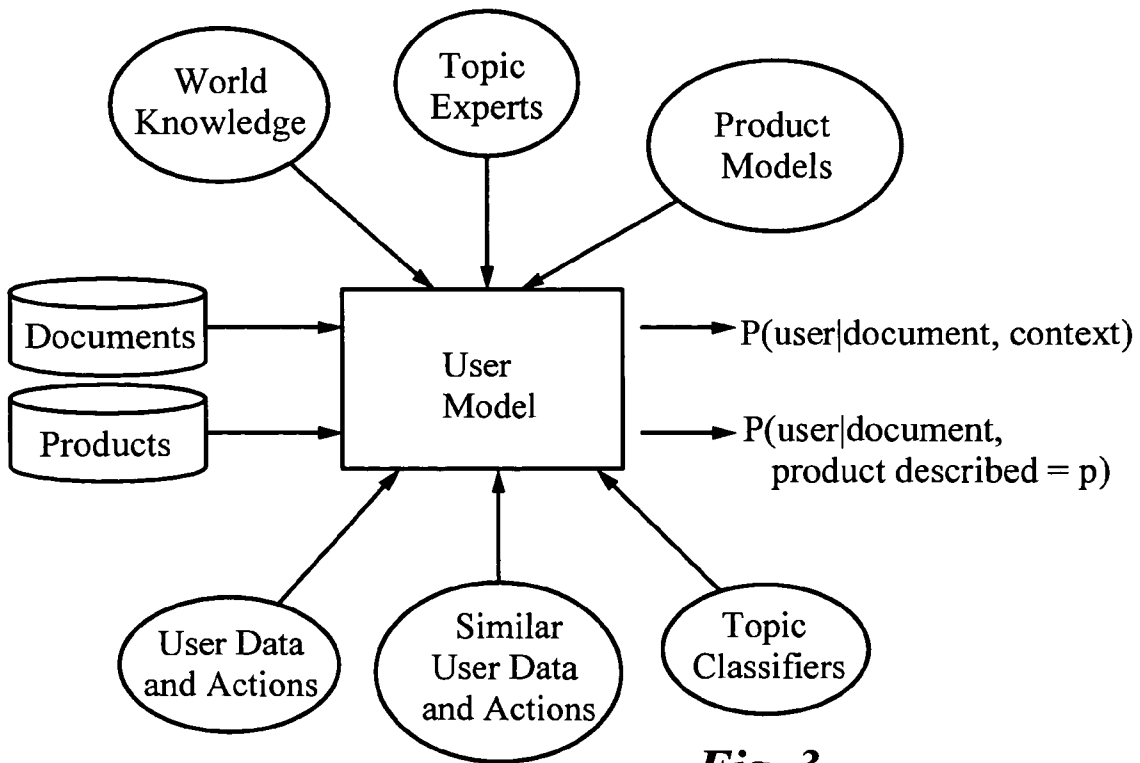


Fig. 3

Informative Word/Phrase List

Word ID	Word Grade	Last Access Time	Number of Accesses
Vegan	0.86	3/6/2000 12:22:41	173
Parasail	0.72	4/15/2000 18:51:27	220

Fig. 4A

Web Site Distribution

Site ID	Site Probability	Last Access Time	Number of Accesses
herring.com	0.61	5/1/2000 19:15:21	152
Java.com	0.43	4/24/2000 3:16:18	460

Fig. 4B

User Topic Distribution

Topic ID	Topic Parent	Topic Probability	Last Access Time	Number of Accesses
Computers	Industries	0.6	12/2/1999 1:21:22	74
Publishing	Industries	0.31	1/2/2000 6:25:31	62

Fig. 4C

User Product Distribution

Product ID	Product Parent	Product Probability	Last Purchase Time	Number of Purchases	Last Access Time	Number of Accesses
3Com Palm 3E	Without Keyboards	0.73	12/16/1999 17:21:21	1	5/2/2000 16:01:21	78
Without Keyboards	Handhelds/PDAs	0.81	12/16/1999 17:21:21	1	3/15/2000 17:21:21	98

Fig. 4D

User Product Feature Distribution

Product ID	Feature ID	Value	Value Probability
Webcams	Interface	PC Card	0.7
Webcams	Interface	Serial	0.2

Fig. 4E

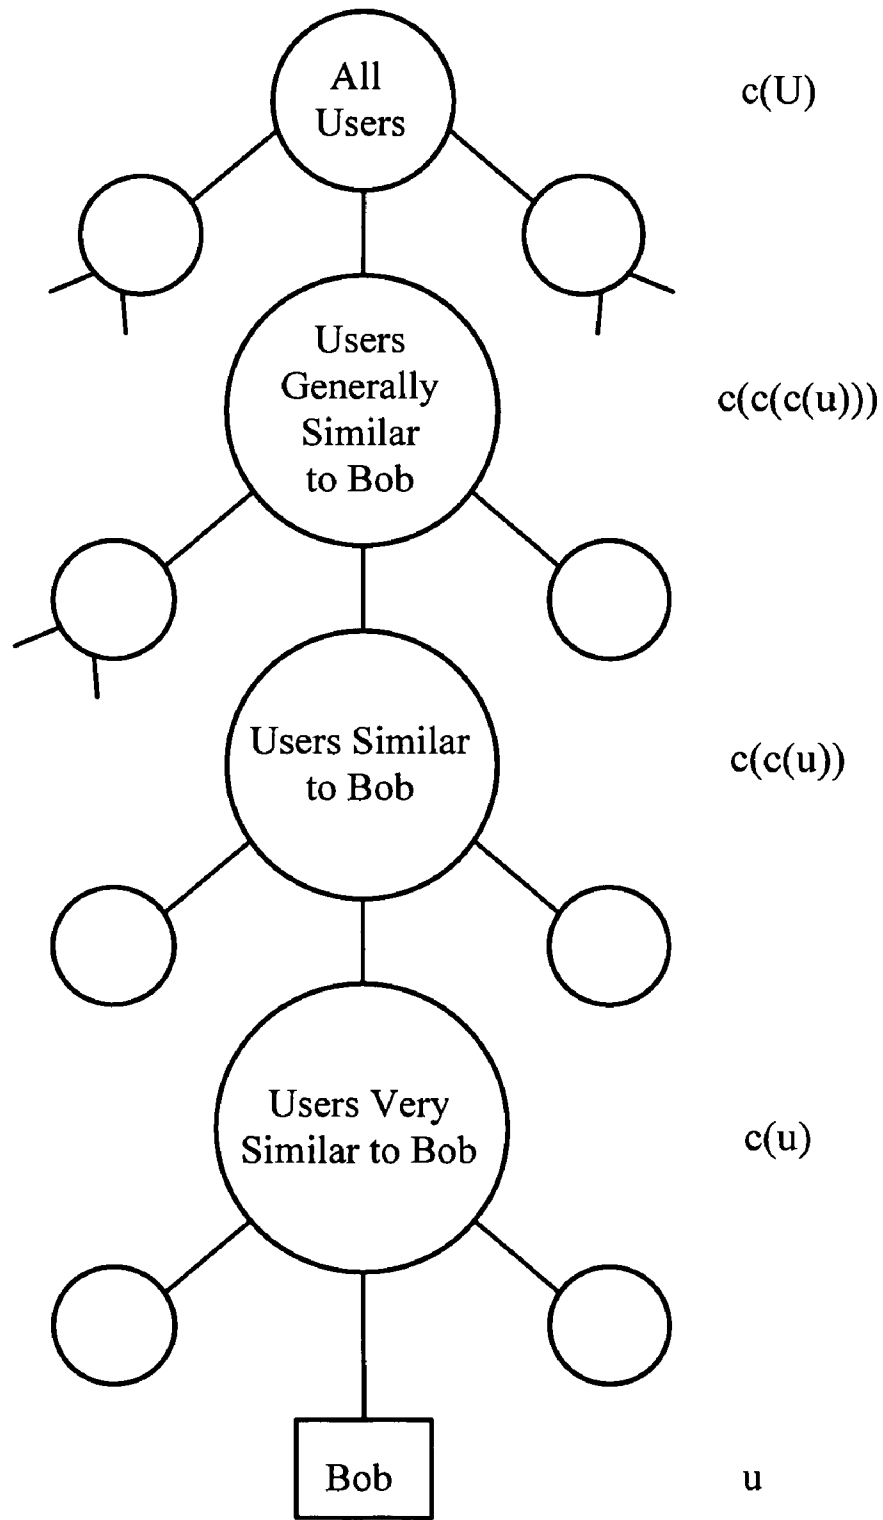


Fig. 5A

User Cluster Tree

Cluster ID	Cluster Parent ID
C123	C3345

Fig. 5B




Fig. 6A

User Fuzzy Cluster Tree

Cluster ID	Cluster Parent ID	Cluster Probability
Bob	C1	0.3
Bob	C2	0.2
Bob	C3	0.1
Bob	C4	0.4
C1	C11	0.2

Fig. 6B




Fig. 7

Topic Tree

Topic ID	Depth Level	Topic Parent ID	Number of Children	Children
Investing	2	Business	3	International, Resources, Socially Responsible
Employment	2	Business	4	Careers, Recruiters, Resumes, Seasonal

Fig. 8

Topic Experts

Topic ID	Topic Parent ID	Cluster 1	Cluster 2
Investing	Business	C112	C113 C114
Employment	Business	C241	C212 C159

Fig. 9

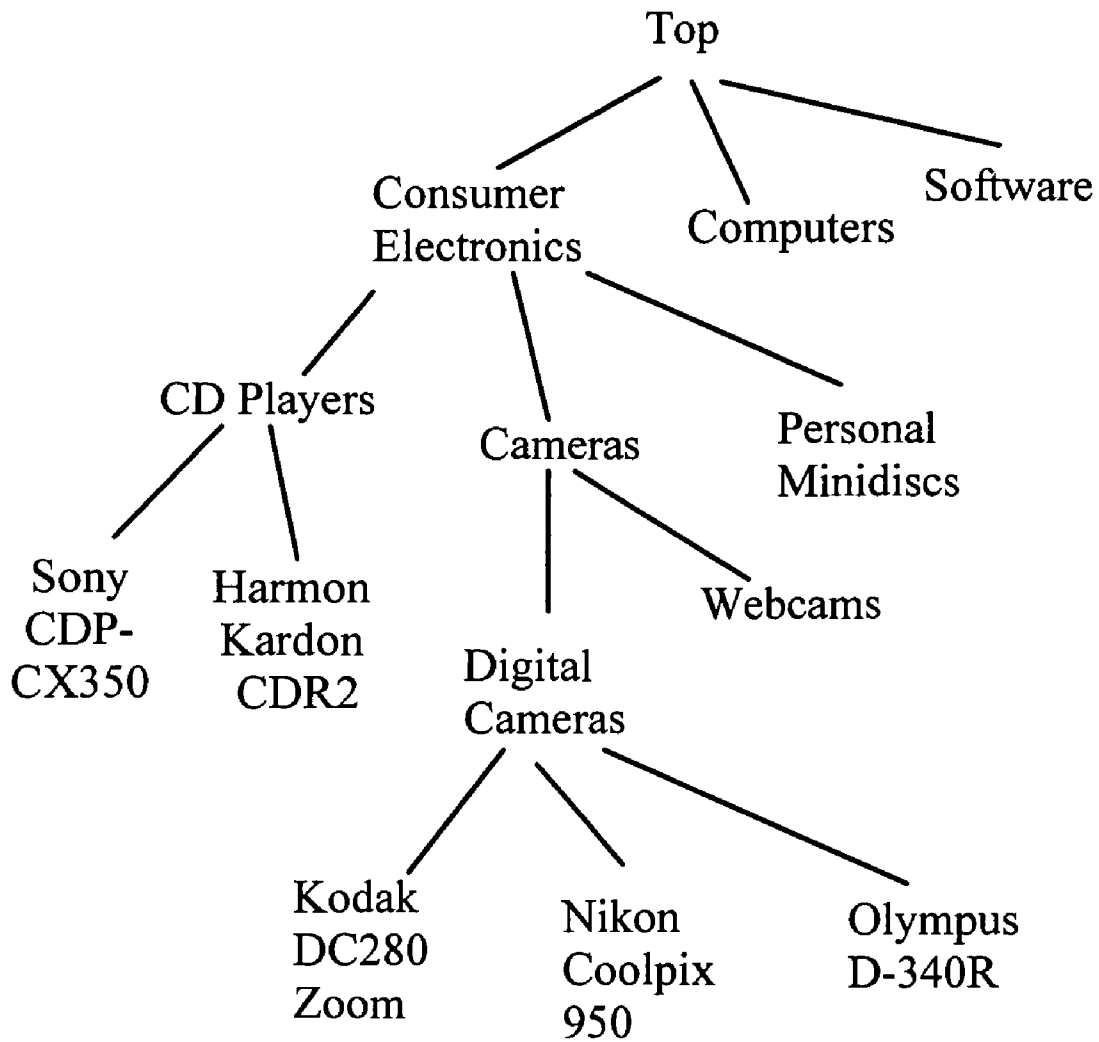


Fig. 10

Product Tree

Product ID	Depth Level	Product Parent ID	Number of Children	Children
Cameras	3	Consumer Electronics	2	Digital Cameras, Webcams
Consumer Electronics	2	Top	3	CD Players, Cameras, Personal Minidiscs

Fig. 11

Product Feature List

Product ID	Feature	Value
Sony CDP-CX350	Brand	Sony
Sony CDP-CX350	CD Capacity	50 Discs or Greater
Sony CDP-CX350	Digital Output	Optical

Fig. 12A

Product Feature Value List

Feature	Value
Digital Output	Coaxial and Optical
Digital Output	Coaxial
Digital Output	Optical
Digital Output	No

Fig. 12B




Fig. 13

User Recently Accessed Butter

Document ID	Access Time	Interaction Type	Context	Degree of Interest
www.herring.com/insider	5/12/2000 14:37:21	Navigation	bookmark access	positive, followed 3 links 12 minutes
www.m-w.com	5/12/2000 15:08:21	Search	query "dictionary"	positive, followed 5 links bookmarked, 21 minutes

Fig. 14

User Site Candidate Table

Site Name	Number of Access	Last Access Time
www.herring.com	157	5/12/2000 14:37:21
www.m-w.com	162	5/12/2000 15:08:21

Fig. 15A

User Word Candidate Table

Word ID	Word Spelling	Word Spelling	Word Grade	Last Access Time
Cytochrome	Cytochrome	Cytocrome	0.67	4/16/2000 7:10:01
Hyperbilirubinemia	Hyperbilirubinemia	Hyperbilirubenema	0.58	4/27/2000 12:18:42

Fig. 15B

User Recently Purchased Products

Product ID	Parent Node	Purchase Time	Purchase Source
Panasonic SL-502	Discmans	5/1/2000 16:01:04	ebyweb.com
Hitachi VM6500A	Camcorders	5/3/2000 18:19:21	supremevideo.com

Fig. 16

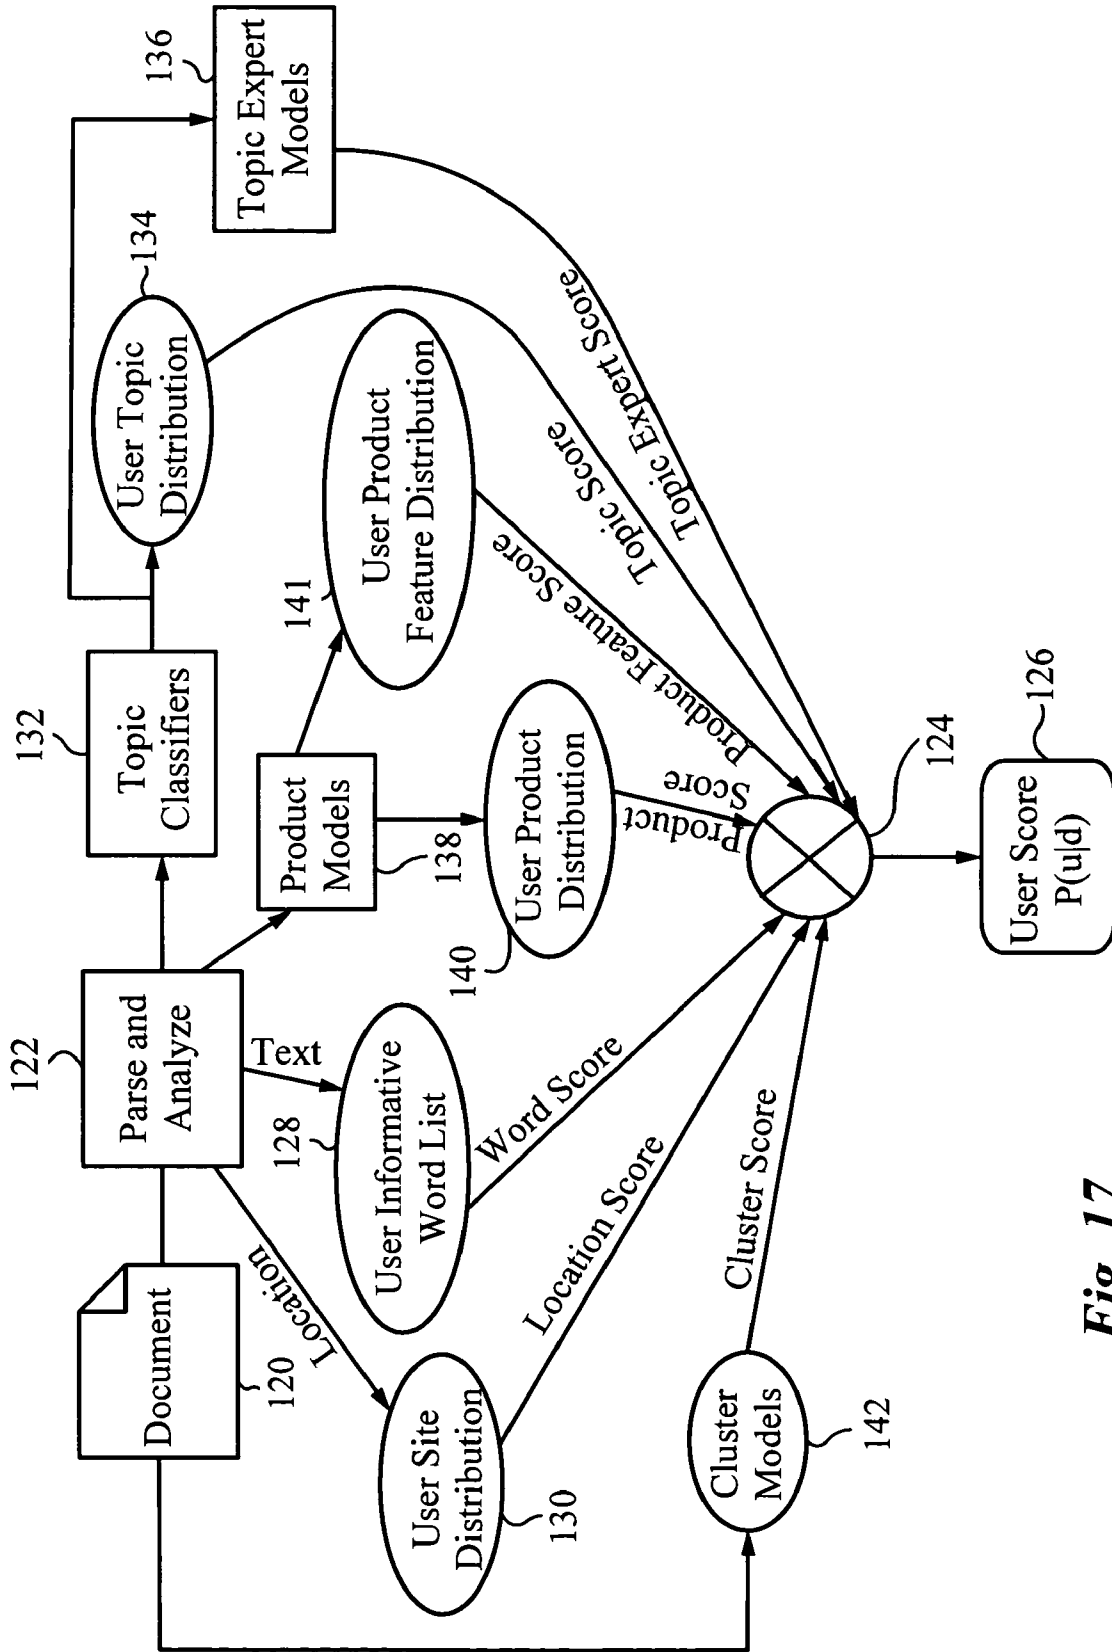


Fig. 17

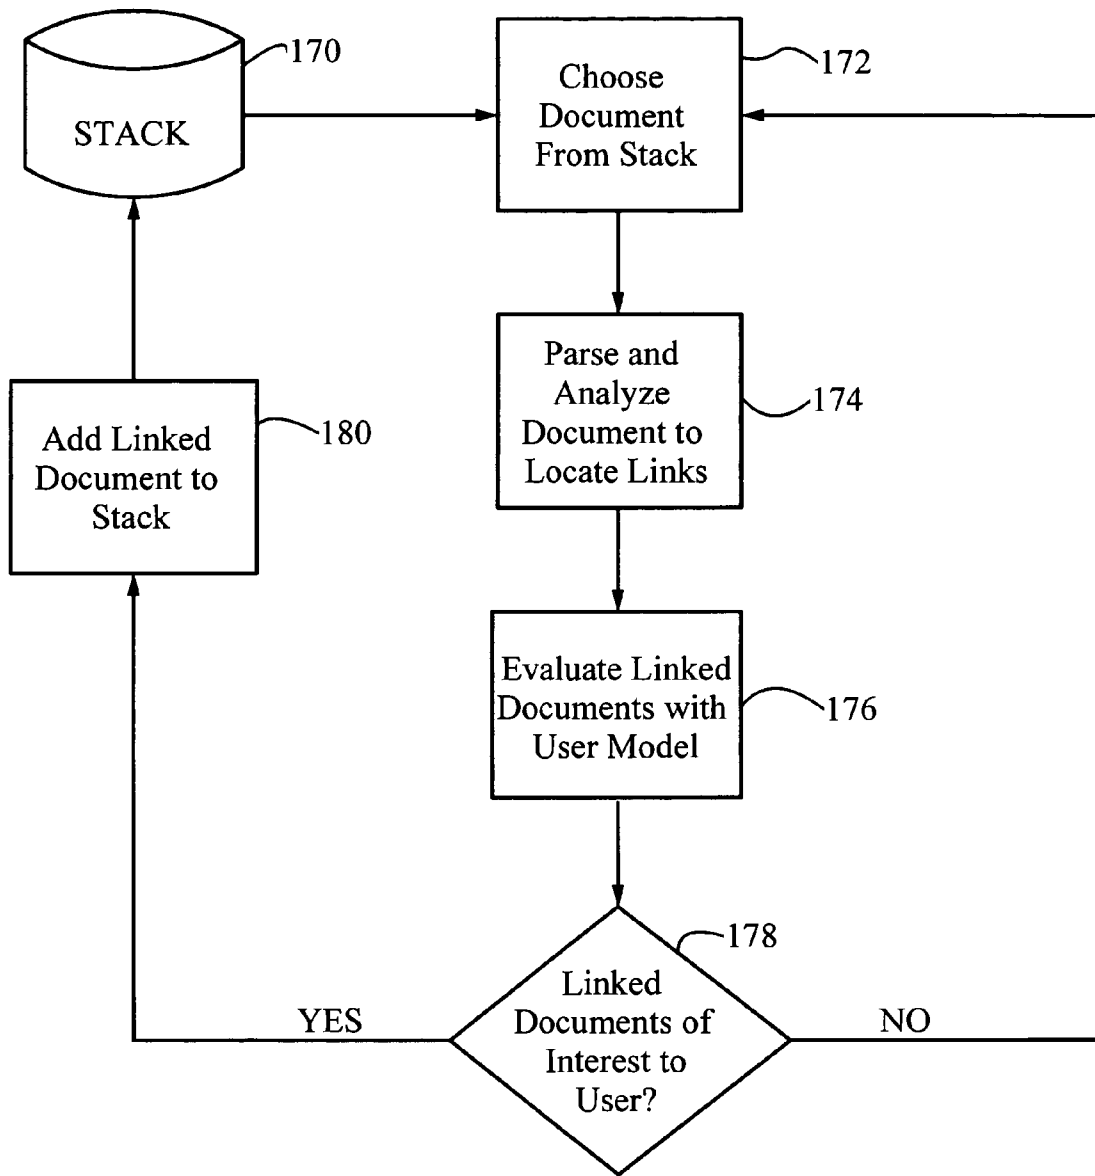


Fig. 18

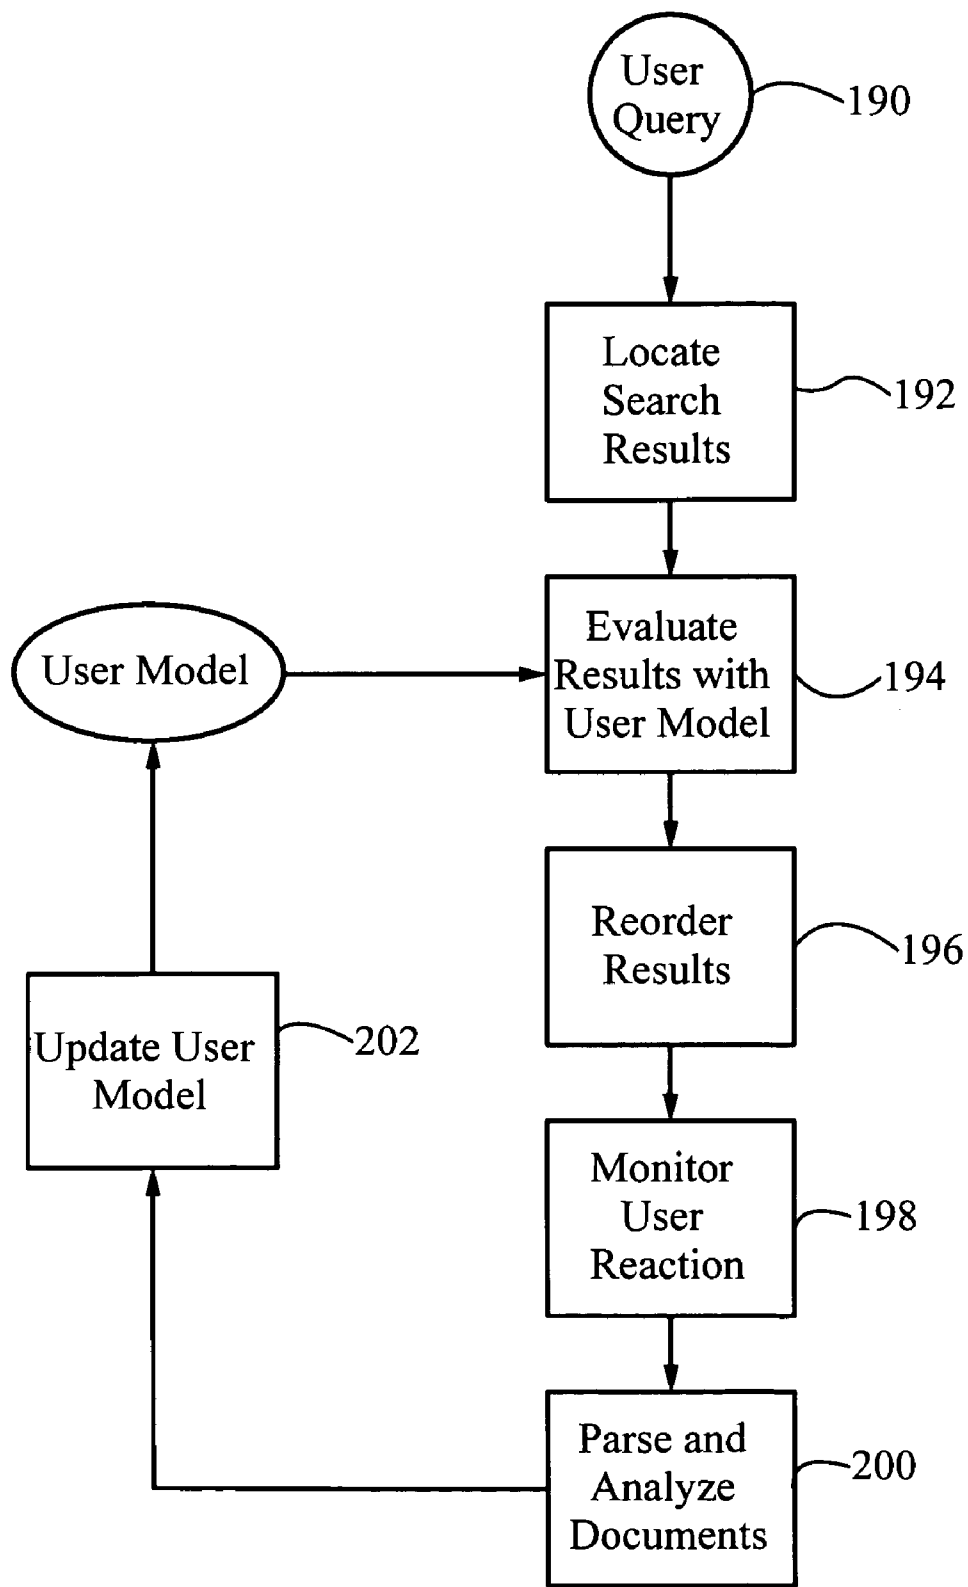


Fig. 19

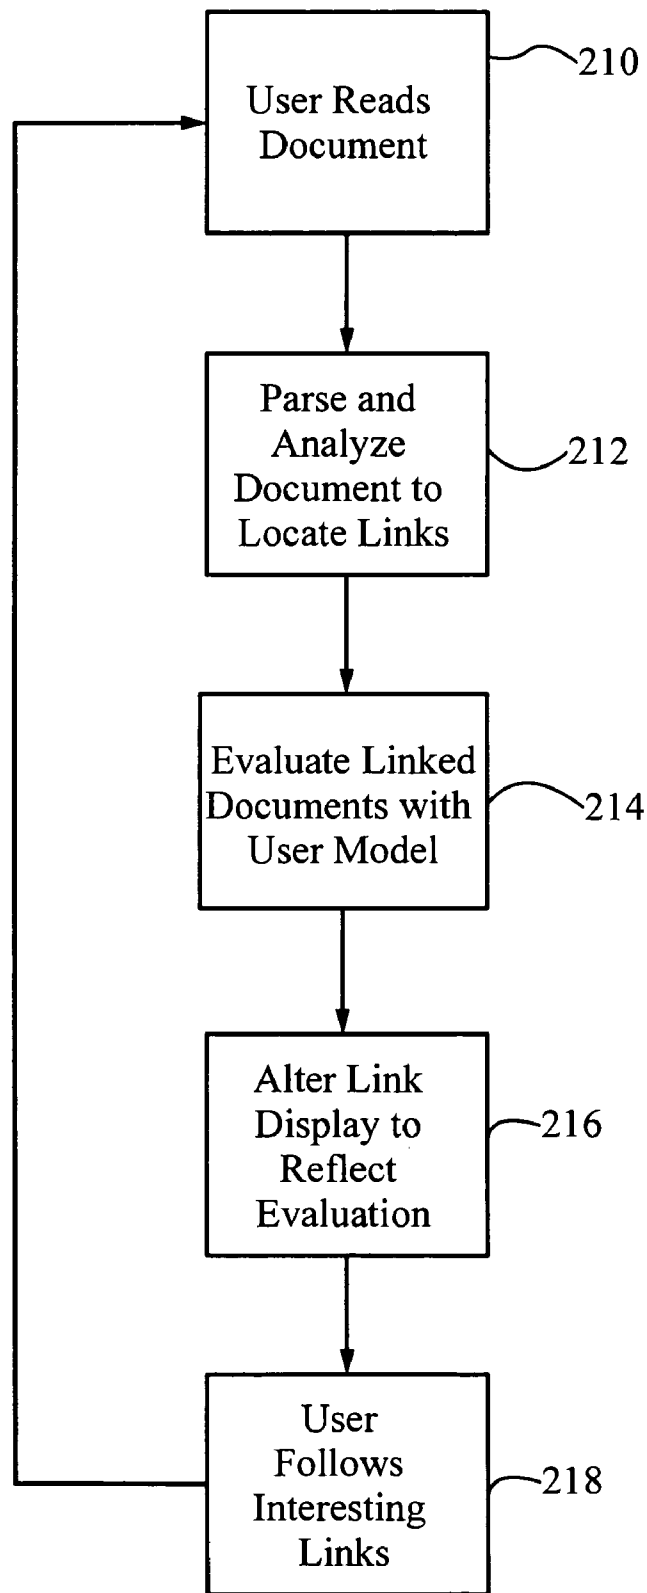


Fig. 20

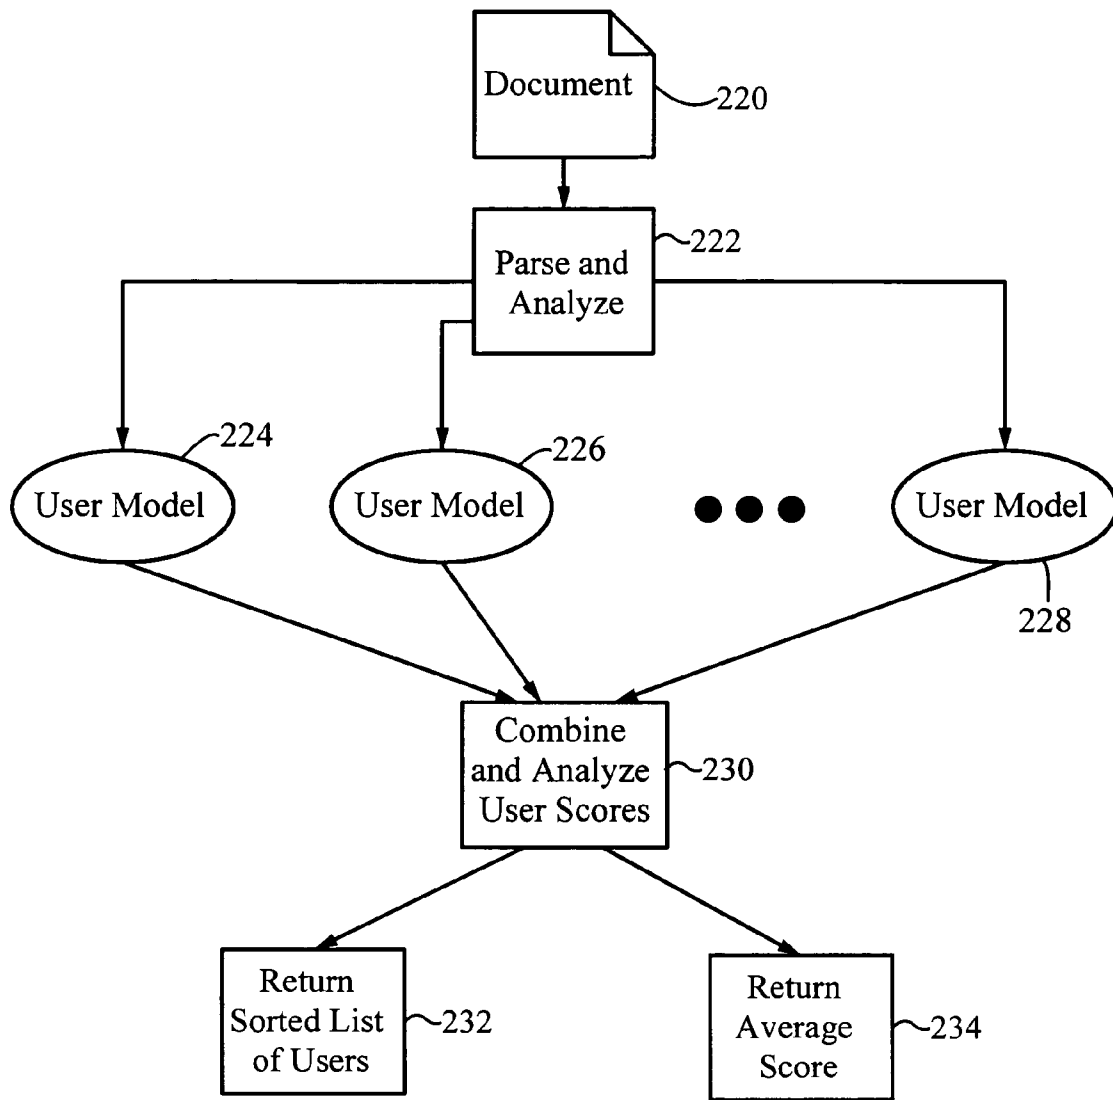


Fig. 21




Fig. 22

AUTOMATIC, PERSONALIZED ONLINE INFORMATION AND PRODUCT SERVICES

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Application No. 60/173,392 filed Dec. 28, 1999, which is herein incorporated by reference.

FIELD OF THE INVENTION

This invention relates generally to methods for personalizing a user's interaction with information in a computer network. More particularly, it relates to methods for predicting user interest in documents and products using a learning machine that is continually updated based on actions of the user and similar users.

BACKGROUND ART

The amount of static and dynamic information available today on the Internet is staggering, and continues to grow exponentially. Users searching for information, news, or products and services are quickly overwhelmed by the volume of information, much of it useless and uninformative. A variety of techniques have been developed to organize, filter, and search for information of interest to a particular user. Broadly, these methods can be divided into information filtering techniques and collaborative filtering techniques.

Information filtering techniques focus on the analysis of item content and the development of a personal user interest profile. In the simplest case, a user is characterized by a set of documents, actions regarding previous documents, and user-defined parameters, and new documents are characterized and compared with the user profile. For example, U.S. Pat. No. 5,933,827, issued to Cole et al., discloses a system for identifying new web pages of interest to a user. The user is characterized simply by a set of categories, and new documents are categorized and compared with the user's profile. U.S. Pat. No. 5,999,975, issued to Kittaka et al., describes an online information providing scheme that characterizes users and documents by a set of attributes, which are compared and updated based on user selection of particular documents. U.S. Pat. No. 6,006,218, issued to Breese et al., discloses a method for retrieving information based on a user's knowledge, in which the probability that a user already knows of a document is calculated based on user-selected parameters or popularity of the document. U.S. Pat. No. 5,754,939, issued to Herz et al., discloses a method for identifying objects of interest to a user based on stored user profiles and target object profiles. Other techniques rate documents using the TFIDF (term frequency, inverse document frequency) measure. The user is represented as a vector of the most informative words in a set of user-associated documents. New documents are parsed to obtain a list of the most informative words, and this list is compared to the user's vector to determine the user's interest in the new document.

Existing information filtering techniques suffer from a number of drawbacks. Information retrieval is typically a two step process, collection followed by filtering; information filtering techniques personalize only the second part of the process. They assume that each user has a personal filter, and that every network document is presented to this filter. This assumption is simply impractical given the current size

and growth of the Internet; the number of web documents is expected to reach several billion in the next few years. Furthermore, the dynamic nature of the documents, e.g., news sites that are continually updated, makes collection of documents to be filtered later a challenging task for any system. User representations are also relatively limited, for example, including only a list of informative words or products or user-chosen parameters, and use only a single mode of interaction to make decisions about different types of documents and interaction modes. In addition, information filtering techniques typically allow for extremely primitive updating of a user profile, if any at all, based on user feedback to recommended documents. As a user's interests change rapidly, most systems are incapable of providing sufficient personalization of a user's experience.

Collaborative filtering methods, in contrast, build databases of user opinions of available items, and then predict a user opinion based on the judgments of similar users. Predictions typically require offline data mining of very large databases to recover association rules and patterns; a significant amount of academic and industrial research is focussed on developing more efficient and accurate data mining techniques. The earliest collaborative filtering systems required explicit ratings by the users, but existing systems are implemented without the user's knowledge by observing user actions. Ratings are inferred from, for example, the amount of time a user spends reading a document or whether a user purchases a particular product. For example, an automatic personalization method is disclosed in B. Mobasher et al., "Automatic Personalization Through Web Usage Mining," Technical Report TR99-010, Department of Computer Science, Depaul University, 1999. Log files of documents requested by users are analyzed to determine usage patterns, and online recommendations of pages to view are supplied to users based on the derived patterns and other pages viewed during the current session.

Recently, a significant number of web sites have begun implementing collaborative filtering techniques, primarily for increasing the number and size of customer purchases. For example, Amazon.com™ has a "Customers Who Bought" feature, which recommends books frequently purchased by customers who also purchased a selected book, or authors whose work is frequently purchased by customers who purchased works of a selected author. This feature uses a simple "shopping basket analysis"; items are considered to be related only if they appear together in a virtual shopping basket. Net Perceptions, an offshoot of the GroupLens project at the University of Minnesota, is a company that provides collaborative filtering to a growing number of web sites based on data mining of server logs and customer transactions, according to predefined customer and product clusters.

Numerous patents disclose improved collaborative filtering systems. A method for item recommendation based on automated collaborative filtering is disclosed in U.S. Pat. No. 6,041,311, issued to Chislenko et al. Similarity factors are maintained for users and for items, allowing predictions based on opinions of other users. In an extension of standard collaborative filtering, item similarity factors allow predictions to be made for a particular item that has not yet been rated, but that is similar to an item that has been rated. A method for determining the best advertisements to show to users is disclosed in U.S. Pat. No. 5,918,014, issued to Robinson. A user is shown a particular advertisement based on the response of a community of similar users to the particular advertisement. New ads are displayed randomly, and the community interest is recorded if enough users click

on the ads. A collaborative filtering system using a belief network is disclosed in U.S. Pat. No. 5,704,317, issued to Heckerman et al., and allows automatic clustering and use of non-numeric attribute values of items. A multi-level mind-pool system for collaborative filtering is disclosed in U.S. Pat. No. 6,029,161, issued to Lang et al. Hierarchies of users are generated containing clusters of users with similar properties.

Collaborative filtering methods also suffer from a number of drawbacks, chief of which is their inability to rate content of an item or incorporate user context. They are based only on user opinions; thus an item that has never been rated cannot be recommended or evaluated. Similarly, obscure items, which are rated by only a few users, are unlikely to be recommended. Furthermore, they require storage of a profile for every item, which is unfeasible when the items are web pages. New items cannot be automatically added into the database. Changing patterns and association rules are not incorporated in real time, since the data mining is performed offline. In addition, user clusters are also static and cannot easily be updated dynamically.

Combinations of information filtering and collaborative filtering techniques have the potential to supply the advantages provided by both methods. For example, U.S. Pat. No. 5,867,799, issued to Lang et al., discloses an information filtering method that incorporates both content-based filtering and collaborative filtering. However, as with content-based methods, the method requires every document to be filtered as it arrives from the network, and also requires storage of a profile of each document. Both of these requirements are unfeasible for realistically large numbers of documents. An extension of this method, described in U.S. Pat. No. 5,983,214, also to Lang et al., observes the actions of users on content profiles representing information entities. Incorporating collaborative information requires that other users have evaluated the exact content profile for which a rating is needed.

In summary, none of the existing prior art methods maintain an adaptive content-based model of a user that changes based on user behavior, allow for real-time updating of the model, operate during the collection stage of information retrieval, can make recommendations for items or documents that have never been evaluated, or model a user based on different modes of interaction.

OBJECTS AND ADVANTAGES

Accordingly, it is a primary object of the present invention to provide a method of personalizing user interaction with network documents that maintains an adaptive content-based profile of the user.

It is another object of the invention to incorporate into the profile user behavior during different modes of interaction with information, thus allowing for cross-fertilization. Learning about the user interests in one mode benefits all other modes.

It is a further object of the invention to provide a method that jointly models the user's information needs and product needs to provide stronger performance in both modes.

It is an additional object of the invention to provide a method that personalizes both the collection and filtering stages of information retrieval to manage efficiently the enormous number of existing web documents.

It is another object of the invention to provide a method for predicting user interest in an item that incorporates the opinions of similar users without requiring storage and maintenance of an item profile.

It is a further object of the invention to provide an information personalization method that models the user as a function independent of any specific representation or data structure, and represents the user interest in a document or product independently of any specific user information need. This approach enables the addition of new knowledge sources into the user model.

It is an additional object of the present invention to provide a method based on Bayesian statistics that updates the user profile based on both negative and positive examples.

It is a further object of the invention to model products by analyzing all relevant knowledge sources, such as press releases, reviews, and articles, so that a product can be recommended even if it has never been purchased or evaluated previously.

SUMMARY

These objects and advantages are attained by a computer-implemented method for providing automatic, personalized information services to a user. User interactions with a computer are transparently monitored while the user is engaged in normal use of the computer, and monitored interactions are used to update user-specific data files that include a set of documents associated with the user. Parameters of a learning machine, which define a User Model specific to the user, are estimated from the user-specific data files. Documents that are of interest and documents that are not of interest to the user are treated distinctly in estimating the parameters. The parameters are used to estimate a probability $P(\text{uld})$ that a document is of interest to the user, and the estimated probability is then used to provide personalized information services to the user.

The probability is estimated by analyzing properties of the document and applying them to the learning machine. Documents of multiple distinct media types of analyzed, and identified properties include: the probability that the document is of interest to users who are interested in particular topics, a topic classifier probability distribution, a product model probability distribution, product feature values extracted from the document, the document author, the document age, a list of documents linked to the document, the document language, number of users who have accessed the document, number of users who have saved the document in a favorite document list, and a list of users previously interested in the document. All properties are independent of the particular user. The product model probability distribution, which indicates the probability that the document refers to particular products, is obtained by applying the document properties to a product model, a learning machine with product parameters characterizing particular products. These product parameters are themselves updated based on the document properties and on the product model probability distribution. Product parameters are initialized from a set of documents associated with each product.

User interactions are monitored during multiple distinct modes of user interaction with network data, including a network searching mode, network navigation mode, network browsing mode, email reading mode, email writing mode, document writing mode, viewing "pushed" information mode, finding expert advice mode, and product purchasing mode. Based on the monitored interactions, parameters of the learning machine are updated. Learning machine parameters define various user-dependent functions of the User Model, including a user topic probability distribution representing interests of the user in various topics, a user

5

product probability distribution representing interests of the user in various products, a user product feature probability distribution representing interests of the user in various features of each of the various products, a web site probability distribution representing interests of the user in various web sites, a cluster probability distribution representing similarity of the user to users in various clusters, and a phrase model probability distribution representing interests of the user in various phrases. Some of the user-dependent functions can be represented as information theory based measures representing mutual information between the user and either phrases, topics, products, features, or web sites. The product and feature distributions can also be used to recommend products to the user.

The User Model is initialized from documents provided by the user, a web browser history file, a web browser bookmarks file, ratings by the user of a set of documents, or previous product purchases made by the user. Alternatively, the User Model may be initialized by selecting a set of predetermined parameters of a prototype user selected by the user. Parameters of the prototype user are updated based on actions of users similar to the prototype user. The User Model can be modified based on User Model modification requests provided by the user. In addition, the user can temporarily use a User Model that is built from a set of predetermined parameters of a profile selected by the user.

Distances between users are calculated to determine similar users, who are clustered into clusters of similar users. Parameters defining the User Model may include the calculated distances between the User Model and User Models of users within the user's cluster. Users may also be clustered based on calculated relative entropy values between User Models of multiple users.

A number of other probabilities can be calculated, such as a posterior probability $P(\text{uld},q)$ that the document is of interest to the user, given a search query submitted by the user. Estimating the posterior probability includes estimating a probability that the query is expressed by the user with an information need contained in the document. In addition, the probability $P(\text{uld},\text{con})$ that the document is of interest to the user during a current interaction session can be calculated. To do so, $P(u,\text{cond})/P(\text{cond})$ is calculated, where con represents a sequence of interactions during the current interaction session or media content currently marked by the user. A posterior probability $P(\text{uld},q,\text{con})$ that the document is of interest to the user, given a search query submitted during a current interaction session, can also be calculated.

A variety of personalized information services are provided using the estimated probabilities. In one application, network documents are crawled and parsed for links, and probable interest of the user in the links is calculated using the learning machine. Links likely to be of interest to the user are followed. In another application, the user identifies a document, and a score derived from the estimated probability is provided to the user. In an additional application, the user is provided with a three-dimensional map indicating user interest in each document of a hyperlinked document collection. In a further application, an expert user is selected from a group of users. The expert user has an expert User Model that indicates a strong interest in a document associated with a particular area of expertise. Another application includes parsing a viewed document for hyperlinks and separately estimating for each hyperlink a probability that the linked document is of interest to the user. In a further application, user interest information derived from the User Model is sent to a third party web server that then customizes its interaction with the user. Finally, a set of users

6

interested in a document is identified, and a range of interests for the identified users is calculated.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 is a schematic diagram of a computer system in which the present invention is implemented.

FIG. 2 is a block diagram of a method of the present invention for providing personalized product and information services to a user.

FIG. 3 is a schematic diagram of knowledge sources used as inputs to the User Model and resulting outputs.

FIGS. 4A–4E illustrate tables that store different components and parameters of the User Model.

FIG. 5A illustrates a cluster tree containing clusters of users similar to a particular user.

FIG. 5B is a table that stores parameters of a user cluster tree.

FIG. 6A illustrates a preferred cluster tree for implementing fuzzy or probabilistic clustering.

FIG. 6B is a table that stores parameters of a user fuzzy cluster tree.

FIG. 7 illustrates a portion of a topic tree.

FIG. 8 is a table that stores nodes of the topic tree of FIG. 7.

FIG. 9 is a table that stores the names of clusters having the most interest in nodes of the topic tree of FIG. 7, used to implement the topic experts model.

FIG. 10 illustrates a portion of a product tree.

FIG. 11 is a table that stores nodes of the product tree of FIG. 10.

FIG. 12A is a table that stores feature values of products of the product tree of FIG. 10.

FIG. 12B is a table that stores potential values of product features associated with intermediate nodes of the product tree of FIG. 10.

FIG. 13 is a schematic diagram of the method of initializing the User Model.

FIG. 14 illustrates the user recently accessed buffer, which records all user interactions with documents.

FIG. 15A is a table for storing sites that are candidates to include in the user site distribution.

FIG. 15B is a table for storing words that are candidates to include in the user word distribution.

FIG. 16 is a table that records all products the user has purchased.

FIG. 17 is a schematic diagram of the method of applying the User Model to new documents to estimate the probability of user interest in the document.

FIG. 18 is a block diagram of the personal crawler application of the present invention.

FIG. 19 is a block diagram of the personal search application of the present invention.

FIG. 20 is a block diagram of the personal navigation application of the present invention.

FIG. 21 is a block diagram of the document barometer application of the present invention.

FIG. 22 is a schematic diagram of the three-dimensional map application of the present invention.

DETAILED DESCRIPTION

Although the following detailed description contains many specifics for the purposes of illustration, anyone of ordinary skill in the art will appreciate that many variations and alterations to the following details are within the scope of the invention. Accordingly, the following preferred

embodiment of the invention is set forth without any loss of generality to, and without imposing limitations upon, the claimed invention.

The present invention, referred to as Personal Web, provides automatic, personalized information and product services to a computer network user. In particular, Personal Web is a user-controlled, web-centric service that creates for each user a personalized perspective and the ability to find and connect with information on the Internet, in computer networks, and from human experts that best matches his or her interests and needs. A computer system **10** implementing Personal Web **12** is illustrated schematically in FIG. 1. Personal Web **12** is stored on a central computer or server **14** on a computer network, in this case the Internet **16**, and interacts with client machines **18, 20, 22, 24, 26** via client-side software. Personal Web **12** may also be stored on more than one central computers or servers that interact over the network. The client-side software may be part of a web browser, such as Netscape Navigator or Microsoft Internet Explorer, configured to interact with Personal Web **12**, or it may be distinct from but interacting with a client browser. Five client machines are illustrated for simplicity, but Personal Web **12** is intended to provide personalized web services for a large number of clients simultaneously.

For all of the typical interactions that a user has with a computer network, such as the world wide web, Personal Web **12** provides a personalized version. Personal Web **12** stores for each user a User Model **13** that is continuously and transparently updated based on the user's interaction with the network, and which allows for personalization of all interaction modes. The User Model represents the user's information and product interests; all information that is presented to the user has been evaluated by the User Model to be of interest to the user. The User Model allows for cross fertilization; that is, information that is learned in one mode of interaction is used to improve performance in all modes of interaction. The User Model is described in detail below.

Five examples of personalized interaction modes provided by the present invention are illustrated in FIG. 1. However, it is to be understood that the present invention provides for personalization of all modes, and that the following examples in no way limit the scope of the present invention. Personal Web is active during all stages of information processing, including collection, retrieval, filtering, routing, and query answering.

Client **18** performs a search using Personal Web **12** by submitting a query and receiving personalized search results. The personal search feature collects, indexes, and filters documents, and responds to the user query, all based on the user profile stored in the User Model **13**. For example, the same query (e.g., "football game this weekend" or "opera") submitted by a teenager in London and an adult venture capitalist in Menlo Park returns different results based on the personality, interests, and demographics of each user. By personalizing the collection phase, the present invention does not require that all network documents be filtered for a particular user, as does the prior art.

Client **20** browses the web aided by Personal Web **12**. In browsing mode, the contents of a web site are customized according to the User Model **13**. Personal Web interacts with a cooperating web site by supplying User Model information, and a web page authored in a dynamic language (e.g., DHTML) is personalized to the user's profile. In navigation mode, a personal navigation aid suggests to the user relevant links within the visited site or outside it given the context, for example, the current web page and previously visited pages, and knowledge of the user profile.

Client **22** illustrates the find-an-expert mode of Personal Web **12**. The user supplies an expert information or product need in the form of a sample web page or text string, and Personal Web **12** locates an expert in the user's company, circle of friends, or outside groups that has the relevant information and expertise, based on the expert's User Model. The located expert not only has the correct information, but presents it in a manner of most interest to the user, for example, focussing on technical rather than business details of a product.

Client **24** uses the personal pushed information mode of Personal Web **12**. Personal Web **12** collects and presents personal information to a user based on the User Model **13**. The pushed information is not limited to a fixed or category or topic, but includes any information of interest to the user. In communities, organizations, or group of users, the pushed information can include automatic routing and delivery of newly created documents that are relevant to the users.

Finally, client **26** illustrates the product recommendation mode of Personal Web **12**. The user submits a query for information about a product type, and Personal Web **12** locates the products and related information that are most relevant to the user, based on the User Model **13**. As described further below, product information is gathered from all available knowledge sources, such as product reviews and press releases, and Personal Web **12** can recommend a product that has never been purchased or rated by any users.

All of the above features of Personal Web **12** are based on a User Model **13** that represents user interests in a document or product independently of any specific user information need, i.e., not related to a specific query. The User Model **13** is a function that is developed and updated using a variety of knowledge sources and that is independent of a specific representation or data structure. The underlying mathematical framework of the modeling and training algorithms discussed below is based on Bayesian statistics, and in particular on the optimization criterion of maximizing posterior probabilities. In this approach, the User Model is updated based on both positive and negative training examples. For example, a search result at the top of the list that is not visited by the user is a negative training example.

The User Model **13**, with its associated representations, is an implementation of a learning machine. As defined in the art, a learning machine contains tunable parameters that are altered based on past experience. Personal Web **12** stores parameters that define a User Model **13** for each user, and the parameters are continually updated based on monitored user interactions while the user is engaged in normal use of a computer. While a specific embodiment of the learning machine is discussed below, it is to be understood that any model that is a learning machine is within the scope of the present invention.

The present invention can be considered to operate in three different modes: initialization, updating or dynamic learning, and application. In the initialization mode, a User Model **13** is developed or trained based in part on a set of user-specific documents. The remaining two modes are illustrated in the block diagram of FIG. 2. While the user is engaged in normal use of a computer, Personal Web **12** operates in the dynamic learning mode to transparently monitor user interactions with data (step **30**) and update the User Model **13** to reflect the user's current interests and needs. This updating is performed by updating a set of user-specific data files in step **32**, and then using the data files to update the parameters of the User Model **13** in step **34**. The user-specific data files include a set of documents

and products associated with the user, and monitored user interactions with data. Finally, Personal Web 12 applies the User Model 13 to unseen documents, which are first analyzed in step 36, to determine the user's interest in the document (step 38), and performs a variety of services based on the predicted user interest (step 40). In response to the services provided, the user performs a series of actions, and these actions are in turn monitored to further update the User Model 13.

The following notation is used in describing the present invention. The user and his or her associated representation are denoted with u , a user query with q , a document with d , a product or service with p , a web site with s , topic with t , and a term, meaning a word or phrase, with w . The term "document" includes not just text, but any type of media, including, but not limited to, hypertext, database, spreadsheet, image, sound, and video. A single document may have one or multiple distinct media types. Accordingly, the set of all possible documents is D , the set of all users and groups is U , the set of all products and services is P , etc. The user information or product need is a subset of D or P . Probability is denoted with P , and a cluster of users or of clusters with c , with which function semantics are used. For example, $c(c(u))$ is the cluster of clusters in which the user u is a member ("the grandfather cluster"). Note that an explicit notation of world knowledge, such as dictionaries, atlases, and other general knowledge sources, which can be used to estimate the various posterior probabilities, is omitted.

A document classifier is a function whose domain is any document, as defined above, and whose range is the continuous interval $[0, 1]$. For example, a document classifier may be a probability that a document d is of interest to a particular user or a group of users. Specific document classifiers of the present invention are obtained using the User Model 13 and Group Model. The User Model 13 represents the user interest in a document independent of any specific user information need. This estimation is unique to each user. In strict mathematical terms, given a user u and a document d , the User Model 13 estimates the probability $P(uld)$. $P(uld)$ is the probability of the event that the user u is interested in the document d , given everything that is known about the document d . This classifier is extended to include $P(uld,con)$, the probability that a user is interested in a given document based on a user's current context, for example, the web pages visited during a current interaction session.

The Group or Cluster Model is a function that represents the interest level of a group of users in a document independently of any specific information need. For example, for the group of users $c(u)$, the mathematical notation of this probability, which is determined by applying the Group Model to a document d , is $P(c(u)d)$.

A schematic diagram of the User Model is shown in FIG. 3, which illustrates the various knowledge sources (in circles) used as input to the User Model. The knowledge sources are used to initialize and update the User Model, so that it can accurately take documents and generate values of user interest in the documents, given the context of the user interaction. Note that some of the knowledge sources are at the individual user level, while others refer to aggregated data from a group of users, while still others are independent of all users. Also illustrated in FIG. 3 is the ability of the User Model to estimate a user interest in a given product, represented mathematically as the interest of a user in a particular document, given that the document describes the product: $P(\text{userdocument, product described}=p)$. As explained further below, the long-term user interest in a

product is one of many probabilities incorporated into the computation of user interest in all documents, but it can also be incorporated into estimation of a current user interest in a product.

Beginning at the bottom left of FIG. 3, User Data and Actions include all user-dependent inputs to the User Model, including user browser documents, user-supplied documents, other user-supplied data, and user actions, such as browsing, searching, shopping, finding experts, and reading news. Data and actions of similar users are also incorporated into the User Model by clustering all users into a tree of clusters. Clustering users allows estimation of user interests based on the interests of users similar to the user. For example, if the user suddenly searches for information in an area that is new to him or her, the User Model borrows characteristics of User Models of users with similar interests. Topic classifiers are used to classify documents automatically into topics according to a predefined topic tree. Similarly, product models determine the product or product categories, if any, referred to by a document. Product models also extract relevant feature of products from product-related documents. The topic experts input provides input of users with a high interest in a particular topic, as measured by their individual User Models. Finally, the User Model incorporates world knowledge sources that are independent of all users, such as databases of company names, yellow pages, thesauri, dictionaries, and atlases.

User Model Representations

Given the inputs shown in FIG. 3, the User Model is a function that may be implemented with any desired data structure and that is not tied to any specific data structure or representation. The following currently preferred embodiment of abstract data structures that represent the User Model 13 is intended to illustrate, but not limit, the User Model of the present invention. Some of the structures hold data and knowledge at the level of individual users, while others store aggregated data for a group or cluster of users. Initialization of the various data structures of the User Model is described in the following section; the description below is of the structures themselves.

User-dependent inputs are represented by components of the User Model shown in FIGS. 4A-4E. These inputs are shown as tables for illustration purposes, but may be any suitable data structure. The user-dependent components include an informative word or phrase list, a web site distribution, a user topic distribution, a user product distribution, and a user product feature distribution. Each of these user-dependent data structures can be thought of as a vector of most informative or most frequent instances, along with a measure representing its importance to the user.

The informative word and phrase list of FIG. 4A contains the most informative words and phrases found in user documents, along with a measure of each informative phrase or word's importance to the user. As used herein, an "informative phrase" includes groups of words that are not contiguous, but that appear together within a window of a predefined number of words. For example, if a user is interested in the 1999 Melissa computer virus, then the informative phrase might include the words "virus," "Melissa," "security," and "IT," all appearing within a window of 50 words. The sentence "The computer virus Melissa changed the security policy of many IT departments" corresponds to this phrase.

In addition to the words and phrases, the list contains the last access time of a document containing each word or phrase and the total number of accessed documents contain-

ing the words. One embodiment of the informative measure is a word probability distribution $P(w|u)$ representing the interest of a user u in a word or phrase w , as measured by the word's frequency in user documents. Preferably, however, the informative measure is not simply a measure of the word frequency in user documents; common words found in many documents, such as "Internet," provide little information about the particular user's interest. Rather, the informative measure should be high for words that do not appear frequently across the entire set of documents, but whose appearance indicates a strong likelihood of the user's interest in a document. A preferred embodiment uses the TFIDF measure, described in Ricardo Baeza-Yates and Berthier Ribeiro-Neto, *Modern Information Retrieval*, Addison Wesley, 1999, in which TF stands for term frequency, and IDF stands for inverse document frequency. Mathematically, if $f_{u,w}$ denotes the frequency of the word w in user u documents, and D_w denotes the number of documents containing the word w , then the importance of a word w to a user u is proportional to the product $f_{u,w} \cdot D_w$.

A more preferred embodiment of the measure of each word's importance uses a mathematically sound and novel implementation based on information theory principles. In particular, the measure used is the mutual information between two random variables representing the user and the word or phrase. Mutual information is a measure of the amount of information one random variable contains about another; a high degree of mutual information between two random variables implies that knowledge of one random variable reduces the uncertainty in the other random variable.

For the present invention, the concept of mutual information is adapted to apply to probability distributions on words and documents. Assume that there is a document in which the user's interest must be ascertained. The following two questions can be asked: Does the phrase p appear in the document?; and Is the document of interest to the user? Intuitively, knowing the answer to one of the questions reduces the uncertainty in answering the other question. That is, if the word w appears in a different frequency in the documents associated with the user u from its frequency in other documents, it helps reduce the uncertainty in determining the interest of user u in the document.

Through the concept of mutual information, information theory provides the mathematical tools to quantify this intuition in a sound way. For a detailed explanation, see T. Cover and J. Thomas, *Elements of Information Theory*, Wiley, 1991. In this embodiment of the informative measure, two indicator variables are defined. I_w has a value of 1 when the word w appears in a web document and 0 when it does not, and I_u has a value of 1 when a web document is of interest to the user u and 0 when it does not. The mutual information between the two random variables I_w and I_u is defined as:

$$I(I_w; I_u) = \sum_{i_w \in I_w} \sum_{i_u \in I_u} P(i_w, i_u) \log_2 \frac{P(i_w, i_u)}{P(i_w)P(i_u)}$$

The probabilities in this formula are computed over a set of documents of interest to the user and a set of documents not of interest to the user. For example, consider a set of 100 documents of interest to the user, and a set of 900 documents not of interest to the user. Then $P(i_u=1)=0.1$, and $P(i_u=0)=0.9$. Assume that in the combined set of 1000 documents, 150 contain the word "Bob." Then $P(i_w=1)=0.15$, and

$P(i_w=0)=0.85$. In addition, assume that "Bob" appears in all 100 of the documents of interest to the user. $P(i_w, i_u)$ has the following four values:

i_u	i_w	$P(i_w, i_u)$
0	0	850/1000
0	1	50/1000
1	0	0/1000
1	1	100/1000

Using the above formula, the mutual information between the user and word Bob is:

$$\begin{aligned} I(I_{Bob}; I_{user}) &= 850/1000 \log [850/1000/(0.85 * 0.9)] + \\ &\quad 50/1000 \log [50/1000/(0.15 * 0.9)] + \\ &\quad 0/1000 \log [0/1000/(0.1 * 0.85)] + \\ &\quad 100/1000 \log [100/1000/(0.15 * 0.1)] \\ &= 0.16. \end{aligned}$$

Mutual information is a preferred measure for selecting the word and phrase list for each user. The chosen words and phrases have the highest mutual information.

The remaining User Model representations are analogously defined using probability distributions or mutual information. The web site distribution of FIG. 4B contains a list of web sites favored by the user along with a measure of the importance of each site. Given the dynamic nature of the Internet, in which individual documents are constantly being added and deleted, a site is defined through the first backslash (after the www). For example, the uniform resource locator (URL) <http://www.herring.com/companies/2000> . . . is considered as www.herring.com. Sites are truncated unless a specific area within a site is considered a separate site; for example, www.cnn.com/health is considered to be a different site than www.cnn.com/us. Such special cases are decided experimentally based on the amount of data available on each site and the principles of data-driven approaches, described in Vladimir S. Cherkassky and Filip M. Mulier, *Learning from Data: Concepts, Theory, and Methods*, in *Adaptive and Learning Systems for Signal Processing, Communications and Control*, Simon Haykin, series editor, Wiley & Sons, March, 1998. Each site has an importance measure, either a discrete probability distribution, $P(stu)$, representing the interest of user u in a web site s , or the mutual information metric defined above, $I(I_s; I_u)$, representing the mutual information between the user u and a site s . The web site distribution also contains the last access time and number of accesses for each site.

FIG. 4C illustrates the user topic distribution, which represents the interests of the user in various topics. The user topic distribution is determined from a hierarchical, user-independent topic model, for example a topic tree such as the Yahoo directory or the Open Directory Project, available at <http://dmoz.org/>. Each entry in the tree has the following form:

Computers\Internet\WWW\Searching the
Web\Directories\Open Directory Project\

where the topic following a backslash is a child node of the topic preceding the backslash. The topic model is discussed in more detail below.

For each node of the topic tree, a probability is defined that specifies the user interest in the topic. Each level of the topic model is treated distinctly. For example, for the top level of the topic model, there is a distribution in which

$$P(t_i|u)+P(t_j|u)=1,$$

where t_i represents the top level of topics and is the same set of topics for each user, e.g., technology, business, health, etc. $P(t_i|u)$ is the sum of the user probabilities on all top level topics. For each topic level, t_i represents specific interests of each user that are not part of any common interest topics, for instance family and friends' home pages. For lower topic levels, every node in the tree is represented in the user topic distribution by a conditional probability distribution. For example, if the Technology node splits into Internet, Communication, and Semiconductors, then the probability distribution is of the form:

$$P(\text{Internet}|u, \text{Technology})+P(\text{Communication}|u, \text{Technology})+P(\text{Semiconductors}|u, \text{Technology})+P(t_i|u, \text{Technology})=1$$

Rather than probabilities, the mutual information metric defined above may be used; $I(I_i; I_u)$ represents the mutual information between the user u and the topic t . An exemplary data structure shown in FIG. 4C for storing the user topic distribution contains, for each topic, the topic parent node, informative measure, last access time of documents classified into the topic, and number of accesses of documents classified into the topic. Note that the User Model contains an entry for every topic in the tree, some of which have a user probability or mutual information of zero.

The user product distribution of FIG. 4D represents the interests of the user in various products, organized in a hierarchical, user-independent structure such as a tree, in which individual products are located at the leaf nodes of the tree. The product taxonomy is described in further detail below. The product taxonomy is similar to the topic tree. Each entry in the tree has the following form:

Consumer Electronics\Cameras\Webcams\3Com Home-Connect\

where a product or product category following a backslash is a child node of a product category preceding the backslash.

For each node of the product model, a probability is defined that specifies the user interest in that particular product or product category. Each level of the product model is treated distinctly. For example, for the top level of the product hierarchy, there is a distribution in which

$$P(p_i|u)=1,$$

where p_i represents the top level of product categories and is the same for each user, e.g., consumer electronics, computers, software, etc. For lower product category levels, every node in the tree is represented in the user product distribution by a conditional probability distribution. For example, if the Cameras node splits into Webcams and Digital Cameras, then the probability distribution is of the form:

$$P(\text{Webcams}|u, \text{Cameras})+P(\text{Digital Cameras}|u, \text{Cameras})=1$$

Rather than probabilities, the mutual information metric defined above may be used. Then $I(I_p; I_u)$ represents the mutual information between the user u and the product or product category p . An exemplary data structure for storing the user product distribution contains, for each product, the product ID, product parent node, user probability, last pur-

chase time of the product, number of product purchases, last access time of documents related to the product, and number of related documents accessed.

For each product or category on which the user has a nonzero probability, the User Model contains a user product feature distribution on the relevant features, as shown in FIG. 4E. Each product category has associated with it a list of features, and the particular values relevant to the user are stored along with a measure of the value's importance, such as a probability $P(f|u, p)$ or mutual information measure $I(I_f; I_u)$. For example, Webcams have a feature Interface with possible values Ethernet (10BaseT), Parallel, PC Card, serial, USB, and TV. Probability values of each feature sum to one; that is,

$$P(\text{Ethernet}|u, \text{Interface, Webcam})+P(\text{Parallel}|u, \text{Interface, Webcam})+P(\text{PC Card}|u, \text{Interface, Webcam})+P(\text{serial}|u, \text{Interface, Webcam})+P(\text{USB}|u, \text{Interface, Webcam})+P(\text{TV}|u, \text{Interface, Webcam})=1.$$

User probability distributions or mutual information measures are stored for each feature value of each node. Note that there is no user feature value distribution at the leaf nodes, since specific products have particular values of each feature.

Finally, user-dependent components of the User Model include clusters of users similar to the user. Users are clustered into groups, forming a cluster tree. One embodiment of a user cluster tree, shown in FIG. 5A, hard classifies users into clusters that are further clustered. Each user is a member of one and only one cluster. For example, Bob is clustered into a cluster $c(u)$, which is further clustered into clusters of clusters, until the top level cluster is reached $c(U)$. The identity of the user's parent cluster and grandfather cluster is stored as shown in FIG. 5B, and information about the parent cluster is used as input into the User Model. As described below, clusters are computed directly from User Models, and thus need not have a predefined semantic underpinning.

Preferably, the User Model does not use hard clustering, but rather uses soft or fuzzy clustering, also known as probabilistic clustering, in which the user belongs to more than one cluster according to a user cluster distribution $P(c(u))$. FIG. 6A illustrates fuzzy clusters in a cluster hierarchy. In this case, Bob belongs to four different clusters according to the probability distribution shown. Thus Bob is most like the members of cluster C4, but still quite similar to members of clusters C1, C2, C3, and C4. Fuzzy clustering is useful for capturing different interests of a user. For example, a user may be a small business owner, a parent of a small child, and also an avid mountain biker, and therefore need information for all three roles. Probabilistic clustering is described in detail in the Ph.D. thesis of Steven J. Nowlan, "Soft Competitive Adaptation: Neural Network Learning Algorithms Based on Fitting Statistical Mixtures," School of Computer Science, Carnegie Mellon University, Pittsburgh, Pa., 1991. A suitable data structure for representing fuzzy clusters is shown in FIG. 6B. Each row stores the cluster or user ID, one parent ID, and the cluster probability, a measure of similarity between the cluster or user and the parent cluster.

Note that all elements of an individual User Model for a user u also apply to a cluster of users $c(u)$. Thus for each cluster, a Group Model is stored containing an informative word list, a site distribution, a topic distribution, a group product distribution, and a group product feature distribution, each with appropriate measures. For example, $P(plc(u))$ represents the interest of a cluster $c(u)$ in various products p .

The user-dependent User Model representations also include a user general information table, which records global information describing the user, such as the User ID, the number of global accesses, the number of accesses within a recent time period, and pointers to all user data structures.

Other knowledge sources of the User Model are independent of the user and all other users. Topic classifiers are used to classify documents into topics according to a predefined topic tree, an example of which is illustrated in FIG. 7. A variety of topic trees are available on the web, such as the Yahoo directory or Open Directory Project (www.dmoz.org). A topic classifier is a model similar to the user model that estimates the probability that a document belongs to a topic. Every node on the topic tree has a stored topic classifier. Thus the set of all topic classifiers computes a probability distribution of all of the documents in the set of documents D among the topic nodes. For example, the topic classifier in the root node in FIG. 7 estimates the posterior probabilities $P(t|d)$, where t represents the topic of document d and is assigned values from the set {Arts, Business, Health, News, Science, Society}. Similarly, the topic classifier for the Business node estimates the posterior probability $P(t|d, \text{Business})$, where t represents the specific topic of the document d within the Business category. Mathematically, this posterior probability is denoted $P(t(d)=\text{Business}|\text{Investing}|(d)=\text{Business}, d)$, which represents the probability that the subtopic of the document d within Business is Investing, given that the topic is Business. The topic tree is stored as shown in FIG. 8, a table containing, for each node, the topic ID, depth level, topic parent ID, number of child nodes, and topic ID of the child nodes.

The topic experts model estimates the probability that a document is of interest to users who are interested in a particular topic, independent of any specific user information need. Each node of the topic tree has, in addition to a topic classifier, a corresponding topic expert function. Note that the topic classifier and topic expert function are independent; two documents can be about investing, but one of high interest to expert users and the other of no interest to expert users. The topic expert model can be considered an evaluation of the quality of information in a given document. The assumption behind the topic experts model is that the degree of interest of a user in a given topic is his or her weight for predicting the quality or general interest level in a document classified within the particular topic. Obviously there are outliers to this assumption, for example, novice users. However, in general and averaged across many users, this measure is a good indicator of a general interest level in a document. For every topic in the tree, a list of the N clusters with the most interest in the topic based on the cluster topic distribution is stored. The cluster topic distribution is similar to the user topic distribution described above, but is averaged over all users in the cluster. An exemplary data structure for storing the topic experts model is shown in FIG. 9.

Finally, a product model is stored for every node of a product taxonomy tree, illustrated in FIG. 10. Examples of product taxonomy trees can be found at www.cnet.com and www.productopia.com, among other locations. In any product taxonomy tree, the leaf nodes, i.e., the bottom nodes of the tree, correspond to particular products, while higher nodes represent product categories. Product models are similar to topic classifiers and User Models, and are used to determine whether a document is relevant to a particular product or product category. Thus a product model contains a list of informative words, topics, and sites. The set of all

product models computes a probability distribution of all of the documents in the set of documents D among the product nodes. For example, the product model in the root node in FIG. 10 estimates the posterior probabilities $P(p|d)$, where p represents the product referred to in document d and is assigned values from the set {Consumer Electronics, Computers, Software}. Similarly, the product model for the Consumer Electronics node estimates the posterior probability $P(p|d, \text{Consumer Electronics})$, where p represents the product category of the document d within the Consumer Electronics category. Mathematically, this posterior probability is denoted $P(p(d)=\text{Consumer Electronics}/\text{CD Players}|(d)=\text{Consumer Electronics}, d)$, which represents the probability that the subproduct category of the document d within Consumer Electronics is CD Players, given that the product category is Consumer Electronics. The product tree is stored as shown in FIG. 11, a table containing, for each node, the topic ID, depth level, topic parent ID, number of child nodes, and topic ID of the child nodes.

Each node of the product tree has an associated product feature list, which contains particular descriptive features relevant to the product or category. Nodes may have associated feature values; leaf nodes, which represent specific products, have values of all relevant product features. Product feature lists are determined by a human with knowledge of the domain. However, feature values may be determined automatically from relevant knowledge sources as explained below.

For example, in the product tree of FIG. 10, CD Players is the parent node of the particular CD players Sony CDP-CX350 and Harman Kardon CDR2. The product category CD Players has the following features: Brand, CD Capacity, Digital Output, Plays Minidisc, and Price Range. Each feature has a finite number of potential feature values; for example, CD Capacity has potential feature values 1 Disc, 1-10 Discs, 10-50 Discs, or 50 Discs or Greater. Individual products, the child nodes of CD Players, have one value of each feature. For example, the Sony CDP-CX350 has a 300 disc capacity, and thus a feature value of 50 Discs or Greater.

Some product features are relevant to multiple product categories. In this case, product features propagate as high up the product tree as possible. For example, digital cameras have the following product features: PC Compatibility, Macintosh Compatibility, Interfaces, Viewfinder Type, and Price Range. Webcams have the following product features: PC Compatibility, Macintosh Compatibility, Interfaces, Maximum Frames per Second, and Price Range. Common features are stored at the highest possible node of the tree; thus features PC Compatibility, Macintosh Compatibility, and Interfaces are stored at the Cameras node. The Digital Cameras node stores only product feature Viewfinder Type, and the Webcams node stores only product feature Maximum Frames per Second.

Note that product feature Price Range is common to CD Players and Cameras, and also Personal Minidiscs, and thus is propagated up the tree and stored at node Consumer Electronics.

Individual products at leaf nodes inherit relevant features from all of their ancestor nodes. For example, Kodak CD280 inherits the feature Viewfinder Type from its parent; PC Compatibility, Macintosh Compatibility, and Interfaces from its grandparent; and Price Range from its great-grandparent. A product feature list is stored as shown in FIG. 12A, and contains, for each product ID, the associated feature and its value. All potential feature values are stored in a product feature value list, as shown in FIG. 12B.

The system also includes a document database that indexes all documents **D**. The document database records, for each document, a document ID, the full location (the URL of the document), a pointer to data extracted from the document, and the last access time of the document by any user. A word database contains statistics of each word or phrase from all user documents. The word database contains the word ID, full word, and word frequency in all documents **D**, used in calculating informative measures for individual users and clusters.

Initialization of User Model

The User Model is initialized offline using characterizations of user behavior and/or a set of documents associated with the user. Each data structure described above is created during initialization. In other words, the relevant parameters of the learning machine are determined during initialization, and then continually updated online during the update mode.

In one embodiment, the user documents for initializing the User Model are identified by the user's web browser. Most browsers contain files that store user information and are used to minimize network access. In Internet Explorer, these files are known as favorites, cache, and history files. Most commercial browsers, such as Netscape Navigator, have equivalent functionality; for example, bookmarks are equivalent to favorites. Users denote frequently-accessed documents as bookmarks, allowing them to be retrieved simply by selection from the list of bookmarks. The bookmarks file includes for each listing its creation time, last modification time, last visit time, and other information. Bookmarks of documents that have changed since the last user access are preferably deleted from the set of user documents. The Internet Temporary folder contains all of the web pages that the user has opened recently (e.g., within the last 30 days). When a user views a web page, it is copied to this folder and recorded in the cache file, which contains the following fields: location (URL), first access time, and last access time (most recent retrieval from cache). Finally, the history file contains links to all pages that the user has opened within a set time period.

Alternatively, the user supplies a set of documents, not included in any browser files, that represent his or her interests. The User Model can also be initialized from information provided directly by the user. Users may fill out forms, answer questions, or play games that ascertain user interests and preferences. The user may also rate his or her interest in a set of documents provided.

User documents are analyzed as shown in FIG. 13 to determine initial parameters for the various functions of the User Model. A similar analysis is used during updating of the User Model. Note that during updating, both documents that are of interest to the user and documents that are not of interest to the user are analyzed and incorporated into the User Model. The process is as follows. In a first step **82**, the format of documents **80** is identified. In step **84**, documents **80** are parsed and separated into text, images and other non-text media **88**, and formatting. Further processing is applied to the text, such as stemming and tokenization to obtain a set of words and phrases **86**, and information extraction. Through information extraction, links **90** to other documents, email addresses, monetary sums, people's names, and company names are obtained. Processing is performed using natural language processing tools such as LinguistX® and keyword extraction tools such as Thing Finder™, both produced by Inxight (www.inxight.com). Further information on processing techniques can be found in Christopher D. Manning and Hinrich Schütze, *Founda-*

tions of Statistical Natural Language Processing, MIT Press, 1999. Additional processing is applied to images and other non-text media **88**. For example, pattern recognition software determines the content of images, and audio or speech recognition software determines the content of audio. Finally, document locations **94** are obtained.

Parsed portions of the documents and extracted information are processed to initialize or update the user representations in the User Model. In step **96**, user informative words or phrases **98** are obtained from document words and phrases **86**. In one embodiment, a frequency distribution is obtained to calculate a TFIDF measure quantifying user interest in words **98**. Alternatively, mutual information is calculated between the two indicator variables I_w and I_u as explained above. The set of informative words **98** contains words with the highest probabilities or mutual information.

In step **100**, the topic classifiers are applied to all extracted information and portions of documents **80** to obtain a probability distribution $P(\text{td})$ for each document on each node of the topic tree. As a result, each node has a set of probabilities, one for each document, which is averaged to obtain an overall topic node probability. The average probabilities become the initial user topic distribution **102**. If desired, mutual information between the two indicator variables I_w and I_u can be determined as explained above.

Similarly, in step **104**, product models are applied to all extracted information from documents **80** to classify documents according to the product taxonomy tree. From user purchase history **105**, additional product probabilities are obtained. Probabilities for each node are combined, weighting purchases and product-related documents appropriately, to obtain a user product distribution **106**. Note that only some of documents **80** contain product-relevant information and are used to determine the user product distribution **106**. Product models return probabilities of zero for documents that are not product related.

The user product feature distribution **108** can be obtained from different sources. If a user has a nonzero probability for a particular product node, then the feature distribution on that node is obtained from its leaf nodes. For example, if one of the user documents was classified into Kodak DC280 and another into Nikon Coolpix 950, then the user product feature distribution for the Digital Cameras node has a probability of 0.5 for the feature values corresponding to each camera. Feature value distributions propagate throughout the user product feature distributions. For example, if the two cameras are in the same price range, \$300-\$400, then the probability of the value \$300-\$400 of the feature Price Range is 1.0, which propagates up to the Consumer Electronics node (assuming that the user has no other product-related documents falling within Consumer Electronics).

Alternatively, product feature value distributions are obtained only from products that the user has purchased, and not from product-related documents in the set of user documents. Relevant feature values are distributed as high up the tree as appropriate. If the user has not purchased a product characterized by a particular feature, then that feature has a zero probability. Alternatively, the user may explicitly specify his or her preferred feature values for each product category in the user product distribution. User-supplied information may also be combined with feature value distributions obtained from documents or purchases.

Document locations **94** are analyzed (step **110**) to obtain the user site distribution **112**. Analysis takes into account the relative frequency of access of the sites within a recent time period, weighted by factors including how recently a site was accessed, whether it was kept in the favorites or

bookmarks file, and the number of different pages from a single site that were accessed. Values of weighting factors are optimized experimentally using jackknifing and cross-validation techniques described in H. Bourlard and N. Morgan, *Connectionist Speech Recognition: A Hybrid Approach*, Kluwer Academic Publishers, 1994.

Note that there is typically overlap among the different representations of the User Model. For example, a news document announcing the release of a new generation of Microsoft servers has relevant words Microsoft and server. In addition, it is categorized within the product taxonomy under Microsoft servers and the topic taxonomy under computer hardware. This document may affect the user's word list, product distribution, and topic distribution.

After the User Models are initialized for all users, cluster membership can be obtained. Clusters contain users with a high degree of similarity of interests and information needs. A large number of clustering algorithms are available; for examples, see K. Fukunaga, *Statistical Pattern Recognition*, Academic Press, 1990. As discussed above, users are preferably soft clustered into more than one cluster. Preferably, the present invention uses an algorithm based on the relative entropy measure from information theory, a measure of the distance between two probability distributions on the same event space, described in T. Cover and J. Thomas, *Elements of Information Theory*, Chapter 2, Wiley, 1991. Clustering is unsupervised. That is, clusters have no inherent semantic significance; while a cluster might contain users with a high interest in mountain biking, the cluster tree has no knowledge of this fact.

In a preferred embodiment, the relative entropy between two User Model distributions on a fixed set of documents D_{sample} is calculated. D_{sample} is chosen as a good representation of the set of all documents D. Distributions of similar users have low relative entropy, and all pairs of users within a cluster have relative entropy below a threshold value. The User Model of each user is applied to the documents to obtain a probability of interest of each user in each document in the set. The relative entropy between two user distributions for a single document is calculated for each document in the set, and then summed across all documents.

The exact mathematical computation of the relative entropy between two users is as follows. An indicator variable $I_{u,d}$ is assigned to 1 when a document d is of interest to a user u and 0 when it is not. For two users u_1 and u_2 and for any document d, the relative entropy between the corresponding distributions is:

$$D(I_{u_1,d} || I_{u_2,d}) = \sum_{i \in I} P(i_{u_1,d}) \log_2 \frac{P(i_{u_1,d})}{P(i_{u_2,d})}$$

For example, if $P(u_1|d)=0.6$ and $P(u_2|d)=0.9$, then

$$D(I_{u_1,d} || I_{u_2,d}) = 0.4 \log(0.4/0.1) + 0.6 \log(0.6/0.9).$$

The relative entropy can be converted to a metric D' that obeys the triangle inequality:

$$D'(I_1 || I_2) = 0.5 * (D(I_1 || I_2) + D(I_2 || I_1)).$$

For any two users u_1 and u_2 , and for each document in D_{sample} , the metric D' is computed between the corresponding indicator variable distributions on the document. The values for all document are summed, and this sum is the

distance metric for clustering users. This distance is defined as:

$$Distance(u_1, u_2) = \sum_{d_j \in D_{sample}} D^2(I_{u_1,d_j} || I_{u_2,d_j}).$$

An alternative clustering algorithm computes the relative entropy between individual user distributions in the User Model, for example, between all informative word lists, site distributions, etc., of each user. The equations are similar to those above, but compute relative entropy based on indicator variables such as $I_{u,w}$, which is assigned a value of 1 when a word w is of interest to a user u. The calculated distances between individual user distributions on words, sites, topics, and products are summed to get an overall user distance. This second algorithm is significantly less computationally costly than the preferred algorithm above; selection of an algorithm depends on available computing resources. In either case, relative entropy can also be computed between a user and cluster of users.

Each cluster has a Group or Cluster Model that is analogous to a User Model. Cluster Models are generated by averaging each component of its members' User Models. When fuzzy clusters are used, components are weighted by a user's probability of membership in the cluster.

In some cases, initialization is performed without any user-specific information. A user may not have a large bookmarks file or cache, or may not want to disclose any personal information. For such users, prototype users are supplied. A user can choose one or a combination of several prototype User Models, such as the technologist, the art lover, and the sports fan. Predetermined parameters of the selected prototype user are used to initialize the User Model. Users can also opt to add only some parameters of a prototype user to his or her existing User Model by choosing the prototype user's distribution of topics, words, sites, etc. Note that prototype users, unlike clusters, are semantically meaningful. That is, prototype users are trained on a set of documents selected to represent a particular interest. For this reason, prototype users are known as "hats," as the user is trying on the hat of a prototype user.

Users can also choose profiles on a temporary basis, for a particular session only. For example, in a search for a birthday present for his or her teenage daughter, a venture capitalist from Menlo Park may be interested in information most probably offered to teenagers, and hence may choose a teenage girl profile for the search session.

User-independent components are also initialized. The topic classifiers are trained using the set of all possible documents D. For example, D may be the documents classified by the Open Directory Project into its topic tree. Topic classifiers are similar to a User Model, but with a unimodal topic distribution function (i.e., a topic model has a topic distribution value of 1 for itself and 0 for all other topic nodes). The set of documents associated with each leaf node of the topic tree is parsed and analyzed as with the user model to obtain an informative word list and site distribution. When a topic classifier is applied to a new document, the document's words and location are compared with the informative components of the topic classifier to obtain $P(t|d)$. This process is further explained below with reference to computation of $P(u|d)$. Preferably, intermediate nodes of the tree do not have associated word list and site distributions. Rather, the measures for the word list and site distribution of child nodes are used as input to the topic classifier

of their parent nodes. For example, the topic classifier for the Business node of the topic tree of FIG. 7 has as its input the score of the site of the document to be classified according to the site distributions of the topic models of its child nodes, Employment, Industries, and Investing. The classifier can be any non-linear classifier such as one obtained by training a Multilayer Perceptron (MLP) using jackknifing and cross-validation techniques, as described in H. Bourlard and N. Morgan, *Connectionist Speech Recognition: A Hybrid Approach*, Kluwer Academic Publishers, 1994. It can be shown that a MLP can be trained to estimate posterior probabilities; for details, see J. Hertz, A. Krogh, R. Palmer, *Introduction to The Theory of Neural Computation*, Addison-Wesley, 1991.

The topic experts model is initialized by locating for every node in the topic tree the N clusters that are of the same depth in the user cluster tree as the user, and that have the highest interest in the topic, based on their cluster topic distribution. The cluster topic distribution $P(tc(u))$ is simply an average of the user topic distribution $P(tu)$ for each user in the cluster. The topic experts model is used to determine the joint probability that a document and the topic under consideration are of interest to any user, $P(t,d)$. Using Bayes' rule, this term can be approximated by considering the users of the N most relevant clusters.

$$P(t, d) = \sum_{i \in N} P(c_i | t, d) P(t | d) P(d)$$

The topic experts model is, therefore, not a distinct model, but rather an ad hoc combination of user and cluster topic distributions and topic models.

Product models are initialized similarly to User Models and topic classifiers. Each leaf node in the product tree of FIG. 10 has an associated set of documents that have been manually classified according to the product taxonomy. These documents are used to train the product model as shown for the User Model in FIG. 13. As a result, each leaf node of the product tree contains a set of informative words, a topic distribution, and a site distribution. Each node also contains a list of features relevant to that product, which is determined manually. From the documents, values of the relevant features are extracted automatically using information extraction techniques to initialize the feature value list for the product. For example, the value of the CD Capacity is extracted from the document. Information extraction is performed on unstructured text, such as HTML documents, semi-structured text, such as XML documents, and structured text, such as database tables. As with the topic model, a nonlinear function such as a Multilayer Perceptron is used to train the product model.

Preferably, as for topic classifiers, intermediate nodes of the product tree do not have associated word lists, site distributions, and topic distributions. Rather, the measures for the word list, site distribution, and topic distribution of child nodes are used as input to the product models of their parent nodes. Alternatively, each parent node may be trained using the union of all documents of its child nodes.

Updating the User Model

The User Model is a dynamic entity that is refined and updated based on all user actions. User interactions with network data are transparently monitored while the user is engaged in normal use of his or her computer. Multiple distinct modes of interaction of the user are monitored,

including network searching, network navigation, network browsing, email reading, email writing, document writing, viewing pushed information, finding expert advice, product information searching, and product purchasing. As a result of the interactions, the set of user documents and the parameters of each user representation in the User Model are modified.

While any nonlinear function may be used in the User Model (e.g., a Multilayer Perceptron), a key feature of the model is that the parameters are updated based on actual user reactions to documents. The difference between the predicted user interest in a document or product and the actual user interest becomes the optimization criterion for training the model.

Through his or her actions, the user creates positive and negative patterns. Positive examples are documents of interest to a user: search results that are visited following a search query, documents saved in the user favorites or bookmarks file, web sites that the user visits independently of search queries, etc. Negative examples are documents that are not of interest to the user, and include search results that are ignored although appear at the top of the search result, deleted bookmarks, and ignored pushed news or email. Conceptually, positive and negative examples can be viewed as additions to and subtractions from the user data and resources.

Information about each document that the user views is stored in a recently accessed buffer for subsequent analysis. The recently accessed buffer includes information about the document itself and information about the user's interaction with the document. One possible implementation of a buffer is illustrated in FIG. 14; however, any suitable data structure may be used. The recently-accessed buffer contains, for each viewed document, a document identifier (e.g., its URL); the access time of the user interaction with the document; the interaction type, such as search or navigation; the context, such as the search query; and the degree of interest, for example, whether it was positive or negative, saved in the bookmarks file, how long the user spent viewing the document, or whether the user followed any links in the document. Additional information is recorded for different modes of interaction with a document as discussed below.

A metric is determined for each document to indicate whether it is a positive, negative or neutral event; this metric can potentially be any grade between 0 and 1, where 0 is a completely negative event, 1 is a completely positive event, and 0.5 is a neutral event. Previous user interactions may be considered in computing the metric; for example, a web site that the user accesses at a frequency greater than a predetermined threshold frequency is a positive example. After each addition to or subtraction from the set of user documents, the document is parsed and analyzed as for the User Model initialization. Extracted information is incorporated into the User Model.

Because the User Model is constantly and dynamically updated, applying the initialization process for each update is inefficient. Preferably, incremental learning techniques are used to update the User Model. Efficient incremental learning and updating techniques provide for incorporating new items into existing statistics, as long as sufficient statistics are recorded. Details about incremental learning can be found in P. Lee, *Bayesian Statistics*, Oxford University Press, 1989.

After a document stored in the recently accessed buffer is parsed, parsed portions are stored in candidate tables. For example, FIGS. 15A and 15B illustrate a user site candidate table and user word candidate table. The user site candidate

table holds sites that are candidates to move into the user site distribution of FIG. 4B. The site candidate table stores the site name, i.e., the URL until the first backslash, except for special cases; the number of site accesses; and the time of last access. The user word candidate table holds the words or phrases that are candidates to move into the user informative word list of FIG. 4A. It contains a word or phrase ID, alternate spellings (or misspellings) of the word, an informative grade, and a time of last access.

Negative examples provide words, sites, and topics that can be used in several ways. The measure of any item obtained from the negative example may be reduced in the user distribution. For example, if the negative example is from a particular site that is in the user site distribution, then the probability or mutual information of that site is decreased. Alternatively, a list of informative negative items may be stored. The negative items are obtained from negative examples and are used to reduce the score of a document containing negative items.

Documents are added to the buffer during all user modes of interaction with the computer. Interaction modes include network searching, network navigation, network browsing, email reading, email writing, document writing, viewing "pushed" information, finding expert advice, and product purchasing. Different types of information are stored in the buffer for different modes. In network searching, search queries are recorded and all search results added to the buffer, along with whether or not a link was followed and access time for viewed search results. In network browsing, the user browses among linked documents, and each document is added to the buffer, along with its interaction time. In email reading mode, each piece of email is considered to be a document and is added to the buffer. The type of interaction with the email item, such as deleting, storing, or forwarding, the sender of the email, and the recipient list are recorded. In email writing mode, each piece of written email is considered a document and added to the buffer. The recipient of the email is recorded. Documents written during document writing mode are added to the buffer. The user's access time with each piece of pushed information and type of interaction, such as saving or forwarding, are recorded. In finding expert advice mode, the user's interest in expert advice is recorded; interest may be measured by the interaction time with an email from an expert, a user's direct rating of the quality of information received, or other suitable measure.

During a product purchasing mode, a similar buffer is created for purchased products, as shown in FIG. 16. All purchased products are used to update the User Model. The user recently purchased products buffer records for each purchase the product ID, parent node in the product tree, purchase time, and purchase source. Purchased products are used to update the user product distribution and user product feature distribution.

If the user feels that the User Model is not an adequate representation of him or her, the user may submit user modification requests. For example, the user may request that specific web sites, topics, or phrases be added to or deleted from the User Model.

User Models for prototype users (hats) are also updated based on actions of similar users. Obviously, it is desirable for prototype User Models to reflect the current state of the representative interest. New web sites appear constantly, and even new informative words appear regularly. For example, technology-related words are introduced and widely adopted quite rapidly; the word list of the Technologist hat should be updated to reflect such changes.

Prototype User Models are updated using actions that are related to the prototype. Actions include documents, user reactions to documents, and product purchases. There are many ways to determine whether an action is relevant to the prototype user. A document that is a positive example for many users (i.e., a followed search result or bookmarked page) and also has a high probability of interest to the prototype user is added to the set of prototype user documents. Actions of users or clusters who are similar to the prototype user, as measured by the relative entropy between individual distributions (words, sites, etc.), are incorporated into the prototype User Model. Additions to the prototype User Model may be weighted by the relative entropy between the user performing the action and the prototype user. Actions of expert users who have a high degree of interest in topics also of interest to the prototype user are incorporated into the prototype User Model.

Note that users who are trying on hats are not able to change the prototype User Model. Their actions affect their own User Models, but not the prototype User Model. Updates to the prototype User Model are based only on actions of users who are not currently trying on hats.

Product models are also continually updated using incremental learning techniques. As described below, the present invention includes crawling network documents and evaluating each document against User Models. Crawled documents are also evaluated by product models. Documents that are relevant to a particular product, as determined by the computed probability $P(\text{pld})$, are used to update its product model. If a document is determined to be relevant, then each component of the product model is updated accordingly. In addition to the parsing and analysis performed for user documents, information extraction techniques are employed to derive feature values that are compared against feature values of the product model, and also incorporated into the feature value list as necessary. New products can be added to the product tree at any time, with characteristic product feature values extracted from all relevant documents. Relevant documents for updating product models include product releases, discussion group entries, product reviews, news articles, or any other type of document.

By employing dynamically updated product models, the present invention, in contrast with prior art systems, provides for deep analysis of all available product information to create a rich representation of products. The interest of a user in a product can therefore be determined even if the product has never been purchased before, or if the product has only been purchased by a very small number of users.

Applying the User Model to Unseen Documents

The User Model is applied to unseen documents to determine the probability that a document is of interest to the user, or the probability that a document is of interest to a user in a particular context. The basic functionality of this determination is then used in the various applications described in subsequent sections to provide personalized information and product services to the user.

The process of estimating user interest in a particular unseen document **120** is illustrated in FIG. 17. This process has the following three steps:

1. Preprocessing the document as for initialization (step **122**).
2. Calculating an individual score for the document for each element of the user representation (e.g., topic distribution, word list).

25

3. Non-linearly combining (124) individual scores into one score 126, the probability that the user is interested in the unseen document, $P(uld)$.

The second step varies for each individual score. From the parsed text, the words of the document 120 are intersected with the words or phrases in the user informative word list 128. For every word or phrase in common, the stored mutual information between the two indicator variables I_u and I_d is summed to obtain the word score. Alternatively, the TFIDF associated with the word are averaged for every common word or phrase. The location score is given by the probability that the document site is of interest to the user, based on the user site distribution 130.

The topic classifiers 132 are applied to document 120 to determine the probability that the document relates to a particular topic, $P(td)$. The user topic score is obtained by computing the relative entropy between the topic distribution $P(td)$ and the user topic distribution 134, $P(tlu)$. After the document has been classified into topics, the topic expert models 136 are applied as described above to determine a score reflecting the interest of users that are experts in the particular topics of this document.

Similarly, the product models 138 are applied to document 120 to determine which products or product categories it describes, $P(pld)$. From the document product distribution, the product score is obtained by computing the relative entropy between the document product distribution and user product distribution 140, $P(plu)$. For each product having a nonzero value of $P(pld)$, its feature values are given by the product model. The user's measures on each of these feature values, found in the user product feature distribution 141, are averaged to obtain a product feature score for each relevant product. Product feature scores are then averaged to obtain an overall product feature score.

The cluster models 142 of clusters to which the user belongs are applied to the document to obtain $P(c(u)d)$. This group model represents the average interests of all users in the cluster. Conceptually, the cluster model is obtained from the union of all the member users' documents and product purchases. Practically, the cluster model is computed from the User Models by averaging the different distributions of the individual User Models, and not from the documents or purchases themselves. Note that in a recursive way, all users have some impact (relative to their similarity to the user under discussion) on the user score, given that $P(c(u)d)$ is estimated using $P(c(c(u))d)$ as a knowledge source, and so on.

Finally, world knowledge (not shown) is an additional knowledge source that represents the interest of an average user in the document based only on a set of predefined factors. World knowledge factors include facts or knowledge about the document, such as links pointing to and from the document or metadata about the document, for example, its author, publisher, time of publication, age, or language. Also included may be the number of users who have accessed the document, saved it in a favorites list, or been previously interested in the document. World knowledge is represented as a probability between 0 and 1.

In step 124, all individual scores are combined to obtain a composite user score 126 for document 120. Step 124 may be performed by training a Multilayer Perceptron using jackknifing and cross-validation techniques, as described in H. Bourlard and N. Morgan, *Connectionist Speech Recognition: A Hybrid Approach*, Kluwer Academic Publishers, 1994. It has been shown in J. Hertz et al., *Introduction to The*

26

Theory of Neural Computation", Addison-Wesley, 1991, that a Multilayer Perceptron can be trained to estimate posterior probabilities.

The context of a user's interaction can be explicitly represented in calculating the user interest in a document. It is not feasible to update the user model after every newly viewed document or search, but the User Model can be updated effectively instantaneously by incorporating the context of user interactions. Context includes content and location of documents viewed during the current interaction session. For example, if the user visits ten consecutive sites pertaining to computer security, then when the User Model estimates the interest of the user in a document about computer security, it is higher than average. The probability of user interest in a document within the current context con is given by:

$$P(u | d, con) = \frac{P(u, con | d)}{P(con | d)}$$

In some applications, individual scores that are combined in step 124 are themselves useful. In particular, the probability that a user is interested in a given product can be used to suggest product purchases to a user. If a user has previously purchased a product, then the User Model contains a distribution on the product's features. If these features propagate far up the product tree, then they can be used to estimate the probability that the user is interested in a different type of product characterized by similar features. For example, if the user purchases a digital camera that is Windows compatible, then the high probability of this compatibility feature value propagates up the tree to a higher node. Clearly, all computer-related purchases for this user should be Windows compatible. Every product that is a descendent of the node to which the value propagated can be rated based on its compatibility, and Windows-compatible products have a higher probability of being of interest to the user.

The long-term interest of a user in products, represented by $P(plu)$, is distinct from the user's immediate interest in a product p , represented as $P(uld, \text{product described}=p)$. The user's immediate interest is the value used to recommend products to a user. Note that $P(plu)$ does not incorporate the user's distribution on feature values. For example, consider the problem of evaluating a user's interest in a particular camera, the Nikon 320. The user has never read any documents describing the Nikon 320, and so $P(\text{Nikon } 320u)=0$. However, the user's feature distribution for the Cameras node indicates high user interest in all of the feature values characterizing the Nikon 320.

When a given product is evaluated by the User Model, the following measures are combined to obtain $P(uld, \text{product described}=p)$: the probabilities of the product and its ancestor nodes from the user product distribution, $P(plu)$; an average of probabilities of each feature value from the user product feature distribution, $P(fl_u, p)$; a probability from the user's clusters' product distributions, $P(fl_c(u), p)$; and an average of probabilities of feature values from the cluster' product feature distributions, $P(fl_c(u), p)$. The overall product score is determined by non-linearly combining all measures. The cluster model is particularly useful if the user does not have a feature value distribution on products in which the user's interest is being estimated.

Applications

The basic function of estimating the probability that a user is interested in a document or product is exploited to provide different types of personalized services to the user. In each type of service, the user's response to the service provided is monitored to obtain positive and negative examples that are used to update the User Model. Example applications are detailed below. However, it is to be understood that all applications employing a trainable User Model as described above are within the scope of the present invention.

Personal Search

In this application, both the collection and filtering steps of searching are personalized. A set of documents of interest to the user is collected, and then used as part of the domain for subsequent searches. The collected documents may also be used as part of the user documents to update the User Model. The collection step, referred to as Personal Crawler, is illustrated schematically in FIG. 18. A stack 170 is initialized with documents of high interest to the user, such as documents in the bookmarks file or documents specified by the user. If necessary, the stack documents may be selected by rating each document in the general document index according to the User Model. The term "stack" refers to a pushdown stack as described in detail in R. Sedgewick, *Algorithms in C++, Parts 1-4*, Addison-Wesley, 1998.

In step 172, the crawler selects a document from the top of the stack to begin crawling. The document is parsed and analyzed (step 174) to identify any links to other documents. If there are links to other documents, each linked document is scored using the User Model (176). If the linked document is of interest to the user (178), i.e., if $P(u|d)$ exceeds a threshold level, then it is added to the stack in step 180, and the crawler continues crawling from the linked document (step 172). If the document is not of interest to the user, then the crawler selects the next document on the stack to continue crawling.

The subsequent searching step is illustrated in FIG. 19. In response to a query 190, a set of search results is located from the set containing all documents D and user documents obtained during personal crawling. The results are evaluated using the User Model (194) and sorted in order of user interest (196), so that the most interesting documents are listed first. The user reaction to each document in the search results is monitored. Monitored reactions include whether or not a document was viewed or ignored and the time spent viewing the document. Documents to which the user responds positively are parsed and analyzed (200) and then used to update the User Model (202) as described above.

The role of the User Model in filtering the search results in step 194 is based on Bayesian statistics and pattern classification theory. According to pattern classification theory, as detailed in R. Duda and P. Hart, *Pattern Classification and Scene Analysis*, Wiley, 1973, the optimal search result is the one with the highest posterior probability. That is, the optimal result is given by:

$$\text{Max}_D P(u | q, d),$$

where $P(u|q,d)$ is the posterior probability of the event that a document d is of interest to a user u having an information need q . This probability can be expressed as:

$$P(u | q, d) = \frac{P(q | d, u)P(u | d)}{P(q | d)}.$$

The term $P(u|d)$ represents the user interest in the document regardless of the current information need, and is calculated using the User Model. The term $P(q|d,u)$ represents the probability that a user u with an information need of d expresses it in the form of a query q . The term $P(q|d)$ represents the probability that an average user with an information need of d expresses it in the form of a query q . One possible implementation of the latter two terms uses the Hidden Markov Model, described in Christopher D. Manning and Hinrich Schütze, *Foundations of Statistical Natural Language Processing*, MIT Press, 1999.

Search results may also be filtered taking into account the context of user interactions, such as content of a recently viewed page or pages. When the context is included, the relevant equation is:

$$P(u | q, d, con) = \frac{P(q | d, u, con)P(u | d, con)}{P(q | d, con)},$$

where $P(u|d,con)$ is as described above.

The Personal Crawler is also used to collect and index documents for product models. Collected documents are parsed and analyzed to update product models, particularly the list of product feature values, which are extracted from collected documents using information extraction techniques.

In general, searches are performed to retrieve all documents from the set of indexed documents that match the search query. Alternatively, searches can be limited to product-related documents, based on either the user's request, the particular search query, or the user's context. For example, a user is interested in purchasing a new bicycle. In one embodiment, the user selects a check-box or other graphical device to indicate that only product-related documents should be retrieved. When the box is not checked, a search query "bicycle" returns sites of bicycle clubs and newsletters. When the box is checked, only documents that have a nonzero product probability ($P(pld)$) on specific products are returned. Such documents include product pages from web sites of bicycle manufacturers, product reviews, and discussion group entries evaluating specific bicycle models.

Alternatively, the search query itself is used to determine the type of pages to return. For example, a query "bicycle" again returns sites of bicycle clubs and newsletters. However, a query "cannondale bicycle" or "cannondale" returns only product-related pages for Cannondale bicycles. Alternatively, the user's context is used to determine the type of pages to return. If the last ten pages viewed by the user are product-related pages discussing Cannondale bicycles, then the query "bicycle" returns product-related pages for all brands of bicycles that are of interest to the user, as determined by the User Model. In all three possible embodiments, within the allowable subset of documents, the entire document is evaluated by the User Model to estimate the probability that the user is interested in the document.

29

Searches may also be performed for products directly, and not for product-related documents. Results are evaluated using only the user product distribution, user product feature distribution, and product and feature distributions of the user's clusters, as explained above. In general, product searches are performed only at the request of the user, for example by selecting a "product search" tab using a mouse or other input device. A user enters a product category and particular feature values, and a list of products that are estimated to be of high interest to the user is returned. The user is returned some form of list of most interesting products. The list may contain only the product name, and may include descriptions, links to relevant documents, images, or any other appropriate information.

Personal Browsing and Navigation

The present invention personalizes browsing and navigation in a variety of different ways. In the personal web sites application, web sites located on third party servers are written in a script language that enables dynamic tailoring of the site to the user interests. Parameters of the User Model are transferred to the site when a user requests a particular page, and only selected content or links are displayed to the user. In one embodiment, the site has different content possibilities, and each possibility is evaluated by the User Model. For example, the CNN home page includes several potential lead articles, and only the one that is most interesting to the user is displayed. In a second embodiment, links on a page are shown only if the page to which they link is of interest to the user. For example, following the lead article on the CNN home page are links to related articles, and only those of interest to the user are shown or highlighted. One single article has a variety of potential related articles; a story on the Microsoft trial, for example, has related articles exploring legal, technical, and financial ramifications, and only those meeting the user's information needs are displayed.

The personal links application is illustrated in FIG. 20. In this application, the hyperlinks in a document being viewed by the user are graphically altered, e.g., in their color, to indicate the degree of interest of the linked documents to the user. As a user views a document (step 210), the document is parsed and analyzed (212) to locate hyperlinks to other documents. The linked documents are located in step 214 (but not shown to the user), and evaluated with the User Model (214) to estimate the user's interest in each of the linked documents. In step 216, the graphical representation of the linked documents is altered in accordance with the score computed with the User Model. For example, the links may be color coded, with red links being most interesting and blue links being least interesting, changed in size, with large links being most interesting, or changed in transparency, with uninteresting links being faded. If the user follows one of the interesting links (218), then the process is repeated for the newly viewed document (210).

The personal related pages application locates pages related to a viewed page. Upon the user's request (e.g., by clicking a button with a mouse pointer), the related pages are displayed. Related pages are selected from the set of user documents collected by the personal crawler. Implementation is similar to that of the personal search application, with

30

the viewed page serving as the query. Thus the relevant equation becomes

$$P(u | page, d) = \frac{P(page | d, u)P(u | d)}{P(page | d)},$$

with $P(page|d,u)$ representing the probability that a user u with an information need of document d expresses it in the form of the viewed page $page$. $P(page|d)$ represents the probability that an average user with an information need of document d expresses it in the form of the viewed page $page$. These terms can be calculated using the Hidden Markov Model.

Alternatively, related pages or sites may be selected according to the cluster model of clusters to which the user belongs. The most likely site navigation from the viewed site, based on the behavior of the cluster members, is displayed to user upon request.

Related pages are particularly useful in satisfying product information needs. For example, if the user is viewing a product page of a specific printer on the manufacturer's web site, clicking the "related pages" button returns pages comparing this printer to other printers, relevant newsgroup discussions, or pages of comparable printers of different manufacturers. All returned related pages have been evaluated by the User Model to be of interest to the user.

Find the Experts

In this application, expert users are located who meet a particular information or product need of the user. Expert users are users whose User Model indicates a high degree of interest in the information need of the user. The information need is expressed as a document or product that the user identifies as representing his or her need. In this context, a document may be a full document, a document excerpt, including paragraphs, phrases, or words, the top result of a search based on a user query, or an email message requesting help with a particular subject. From the pool of potential experts, User Models are applied to the document or product, and users whose probability of interest in the document or product exceeds a threshold level are considered expert users.

The pool of experts is specified either by the user or in the system. For example, the pool may include all company employees or users who have previously agreed to help and advise other users. When users request expert advice about a particular product, the expert may be chosen from the product manufacturer or from users who have previously purchased the product, or from users participating in discussion groups about the product.

A protocol for linking users and identified experts is determined. For example, the expert receives an email message requesting that he or she contact the user in need of assistance. Alternatively, all user needs are organized in a taxonomy of advice topics, and an expert searches for requests associated with his or her topic of expertise.

Personal News

This application, also known as personal pushed information, uses the personal crawler illustrated in FIG. 18. From all documents collected within a recent time period by the user's crawler or user's clusters' crawlers, the most interesting ones are chosen according to the User Model. Collection sources may also be documents obtained from news wires of actions of other users. Documents are sent to the user in any suitable manner. For example, users receive

email messages containing URLs of interesting pages, or links are displayed on a personal web page that the user visits.

Personalization Assistant

Using the User Model, the Personalization Assistant can transform any services available on the web into personalized services, such as shopping assistants, chatting browsers, or matchmaking assistants.

Document Barometer

The document barometer, or Page-O-Meter, application, illustrated in FIG. 21, finds the average interest of a large group of users in a document. The barometer can be used by third parties, such as marketing or public relations groups, to analyze the interest of user groups in sets of documents, advertising, or sites, and then modify the documents or target advertising at particular user groups. The application can instead report a score for a single user's interest in a document, allowing the user to determine whether the system is properly evaluating his or her interest. If not, the user can make user modification requests for individual elements of the User Model. From individual and average scores, the application determines a specific user or users interested in the document.

Referring to FIG. 21, a document 220 is parsed and analyzed (222) and then evaluated according to a set of N User Models 224 and 226 through 228. N includes any number greater than or equal to one. The resulting scores from all User Models are combined and analyzed in step 230. In one embodiment, the analysis locates users having maximum interest in document 220, or interest above a threshold level, and returns a sorted list of interested users (232). Alternatively, an average score for document 220 is calculated and returned (234). The average score may be for all users or for users whose interest exceeds a threshold interest level. The range of interest levels among all users in the group may also be reported.

An analogous product barometer calculates user interest in a product. The product barometer computes a score for an individual user or group of users, or identifies users having an interest in a product that exceeds a threshold level. Third party organizations use the product barometer to target marketing efforts to users who are highly likely to be interested in particular products.

3D Map

FIG. 22 illustrates a three-dimensional (3D) map 240 of the present invention, in which rectangles represent documents and lines represent hyperlinks between documents. A user provides a set of hyperlinked documents, and each document is scored according to the User Model. An image of 3D map 240 is returned to the user. 3D map 240 contains, for each document, a score reflecting the probability of interest of the user in the document.

Product Recommendations

A user's online shopping experience can be personalized by making use of the user's overall product score described above, $P(uld, \text{product described}=p)$. Products that are of high interest to the user are suggested to him or her for purchase. When a user requests information for a specific product or purchases a product, related products are suggested (up-sell). Related product categories are predetermined by a human, but individual products within related categories are evaluated by the User Model before being suggested to the user. The related products are given to the user in a list that may contain images, hyperlinks to documents, or any other suitable information. For example, when a user purchases a

server, a list of relevant backup tapes are suggested to him or her for purchase. Suggested products may have feature values that are known to be of interest to the user, or may have been purchased by other members of the user's cluster who also purchased the server. Related product suggestions may be made at any time, not only when a user purchases or requests information about a particular product. Suggested products may be related to any previously purchased products.

Similarly, competing or comparable products are suggested to the user (cross-sell). When the user browses pages of a particular product, or begins to purchase a product, products within the same product category are evaluated to estimate the user's interest in them. Products that are highly interesting to the user are recommended. The user might intend to purchase one product, but be shown products that are more useful or interesting to him or her.

It will be clear to one skilled in the art that the above embodiments may be altered in many ways without departing from the scope of the invention. Accordingly, the scope of the invention should be determined by the following claims and their legal equivalents.

What is claimed is:

1. A computer-implemented method for providing automatic, personalized information services to a user u, the method comprising:

- a) transparently monitoring user interactions with data while the user is engaged in normal use of a computer;
- b) updating user-specific data files, wherein the user-specific data files comprise the monitored user interactions with the data and a set of documents associated with the user;
- c) estimating parameters of a learning machine, wherein the parameters define a User Model specific to the user and wherein the parameters are estimated in part from the user-specific data files;
- d) analyzing a document d to identify properties of the document;
- e) estimating a probability $P(uld)$ that an unseen document d is of interest to the user u, wherein the probability $P(uld)$ is estimated by applying the identified properties of the document to the learning machine having the parameters defined by the User Model; and
- f) using the estimated probability to provide automatic, personalized information services to the user.

2. The method of claim 1 wherein the user-specific data files include documents of interest to the user u and documents that are not of interest to the user u, and wherein estimating the parameters comprises distinct treatment of the documents of interest and the documents that are not of interest.

3. The method of claim 1 wherein analyzing the document d provides for the analysis of documents having multiple distinct media types.

4. The method of claim 1 wherein transparently monitoring user interactions with data comprises monitoring multiple distinct modes of user interaction with network data.

5. The method of claim 4 wherein the multiple distinct modes of user interaction comprise a mode selected from the group consisting of a network searching mode, a network navigation mode, a network browsing mode, an email reading mode, an email writing mode, a document writing mode, a viewing "pushed" information mode, a finding expert advice mode, and a product purchasing mode.

6. The method of claim 1 further comprising crawling network documents, wherein the crawling comprises parsing crawled documents for links, calculating probable user inter-

est in the parsed links using the learning machine, and preferentially following links likely to be of interest to the user.

7. The method of claim 1 wherein the identified properties of the document d comprise a user u-independent property selected from the group consisting of:

- a) a probability $P(t,d)$ that the document d is of interest to users interested in a topic t;
- b) a topic classifier discrete probability distribution $P(td)$;
- c) a product model discrete probability distribution $P(pld)$;
- d) product feature values extracted from the document d;
- e) an author of the document d;
- f) an age of the document d;
- g) a list of documents linked to the document d;
- h) a language of the document d;
- i) a number of users who have accessed the document d;
- j) a number of users who have saved the document d in a favorite document list; and
- k) a list of users previously interested in the document d.

8. The method of claim 1 wherein the parameters of the learning machine define a user u-dependent function selected from the group consisting of:

- a) a user topic probability distribution $P(tu)$ representing interests of the user u in various topics t;
- b) a user product probability distribution $P(plu)$ representing interests of the user u in various products p;
- c) a user product feature probability distribution $P(fl_u,p)$ representing interests of the user u in various features f of each of the various products p;
- d) a web site probability distribution $P(slu)$ representing interests of the user u in various web sites s;
- e) a cluster probability distribution $P(c(u)u)$ representing similarity of the user u to users in various clusters c(u);
- f) a phrase model probability distribution $P(wu)$ representing interests of the user u in various phrases w;
- g) an information theory based measure $I(I_w; I_u)$ representing mutual information between various phrases w and the user u;
- h) an information theory based measure $I(I_t; I_u)$ representing mutual information between various topics t and the user u;
- i) an information theory based measure $I(I_s; I_u)$ representing mutual information between various web sites s and the user u;
- j) an information theory based measure $I(I_p; I_u)$ representing mutual information between various products p and the user u; and
- k) an information theory based measure $I(I_f; I_u)$ representing mutual information between various features f of each of the various products p and the user u.

9. The method of claim 1 wherein the parameters of the learning machine define:

- a) a user product probability distribution $P(plu)$ representing interests of the user u in various products p; and
- b) a user product feature probability distribution $P(fl_u,p)$ representing interests of the user u in various features f of each of the various products p; and wherein the method further comprises estimating a probability $P(uld, \text{product described}=p)$ that a document d that describes a product p is of interest to the user u, wherein the probability is estimated in part from the user product probability distribution and the user product feature probability distribution.

10. The method of claim 9 further comprising recommending products to the user based on the probability $P(uld, \text{product described}=p)$.

11. The method of claim 1 further comprising estimating a posterior probability $P(uld,q)$ that the document d is of interest to the user u, given a query q submitted by the user.

12. The method of claim 11 wherein estimating the posterior probability comprises estimating a probability $P(qd,u)$ that the query q is expressed by the user u with an information need in the document d.

13. The method of claim 1 further comprising applying the identified properties of the document d to a learning machine having product parameters characterizing a product p to estimate a probability $P(pld)$ that the document d refers to the product p.

14. The method of claim 13 further comprising updating the product parameters based on the identified properties of the document d and the estimated probability $P(pld)$.

15. The method of claim 13 further comprising initializing the product parameters based on a set of documents associated with the product p.

16. The method of claim 1 further comprising clustering multiple users into clusters of similar users, wherein the clustering comprises calculating distances between User Models, and selecting similar users based on the calculated distances between User Models.

17. The method of claim 1 further comprising calculating relative entropy values between User Models of multiple users, and clustering together users based on the calculated relative entropy values.

18. The method of claim 1 wherein the parameters defining the User Model comprise calculated distances between the User Model and User Models of users similar to the user.

19. The method of claim 1 further comprising selecting in a group of users an expert user in an area of expertise, wherein selecting the expert user comprises finding an expert User Model among User Models of the group of users, such that the expert User Model indicates a strong interest of the expert user in a document associated with the area of expertise.

20. The method of claim 1 further comprising parsing the document d for hyperlinks, and separately estimating for each of the hyperlinks a probability that the hyperlink is of interest to the user u.

21. The method of claim 1 further comprising sending to a third party web server user interest information derived from the User Model, whereby the third party web server may customize its interaction with the user.

22. The method of claim 1 wherein the monitored user interactions include a sequence of interaction times.

23. The method of claim 1 further comprising initializing the User Model using information selected from the group consisting of a set of documents provided by the user, a web browser history file associated with the user, a web browser bookmarks file associated with the user, ratings by the user of a set of documents, and previous product purchases made by the user.

24. The method of claim 1 further comprising modifying the User Model based on User Model modification requests provided by the user.

25. The method of claim 1 further comprising providing to the user a score for a document identified by the user, wherein the score is derived from the estimated probability.

26. The method of claim 1 further comprising providing to the user a 3D map of a hyper linked document collection, wherein the 3D map indicates a user interest in each document.

27. The method of claim 1 further comprising temporarily using a User Model that is built from a set of predetermined parameters of a profile selected by the user.

35

28. The method of claim 1 further comprising initializing the User Model by selecting a set of predetermined parameters of a prototype user selected by the user.

29. The method of claim 28 further comprising updating the predetermined parameters of the prototype user based on actions of users similar to the prototype user.

30. The method of claim 1 further comprising identifying a set of users interested in the document d.

31. The method of claim 30 further comprising calculating a range of interests in the document d for the identified set of users.

32. A program storage device accessible by a central computer, tangibly embodying a program of instructions executable by the central computer to perform method steps for providing automatic, personalized information services to a user u, the method steps comprising:

- a) transparently monitoring user interactions with data while the user is engaged in normal use of a client computer in communication with the central computer;
- b) updating user-specific data files, wherein the user-specific data files comprise the monitored user interactions with the data and a set of documents associated with the user;
- c) estimating parameters of a learning machine, wherein the parameters define a User Model specific to the user and wherein the parameters are estimated in part from the user-specific data files;
- d) analyzing a document d to identify properties of the document;
- e) estimating a probability $P(uld)$ that an unseen document d is of interest to the user u, wherein the probability $P(uld)$ is estimated by applying the identified properties of the document to the learning machine having the parameters defined by the User Model; and
- f) using the estimated probability to provide automatic, personalized information services to the user.

33. The program storage device of claim 32 wherein the user-specific data files include documents of interest to the user u and documents that are not of interest to the user u, and wherein estimating the parameters comprises distinct treatment of the documents of interest and the documents that are not of interest.

34. The program storage device of claim 32 wherein analyzing the document d provides for the analysis of documents having multiple distinct media types.

35. The program storage device of claim 32 wherein transparently monitoring user interactions with data comprises monitoring multiple distinct modes of user interaction with network data.

36. The program storage device of claim 35 wherein the multiple distinct modes of user interaction comprise a mode selected from the group consisting of a network searching mode, a network navigation mode, a network browsing mode, an email reading mode, an email writing mode, a document writing mode, a viewing "pushed" information mode, a finding expert advice mode, and a product purchasing mode.

37. The program storage device of claim 32 wherein the method steps further comprise crawling network documents, wherein the crawling comprises parsing crawled documents for links, calculating probable user interest in the parsed links using the learning machine, and preferentially following links likely to be of interest to the user.

38. The program storage device of claim 32 wherein the identified properties of the document d comprise a user

36

u-independent property selected from the group consisting of:

- a) a probability $P(t,d)$ that the document d is of interest to users interested in a topic t;
- b) a topic classifier discrete probability distribution $P(t|d)$;
- c) a product model discrete probability distribution $P(p|d)$;
- d) product feature values extracted from the document d;
- e) an author of the document d;
- f) an age of the document d;
- g) a list of documents linked to the document d;
- h) a language of the document d;
- i) a number of users who have accessed the document d;
- j) a number of users who have saved the document d in a favorite document list; and
- k) a list of users previously interested in the document d.

39. The program storage device of claim 32 wherein the parameters of the learning machine define a user u-dependent function selected from the group consisting of:

- a) a user topic probability distribution $P(t|u)$ representing interests of the user u in various topics t;
- b) a user product probability distribution $P(p|u)$ representing interests of the user u in various products p;
- c) a user product feature probability distribution $P(f|u,p)$ representing interests of the user u in various features f of each of the various products p;
- d) a web site probability distribution $P(s|u)$ representing interests of the user u in various web sites s;
- e) a cluster probability distribution $P(c(u)|u)$ representing similarity of the user u to users in various clusters c(u);
- f) a phrase model probability distribution $P(w|u)$ representing interests of the user u in various phrases w;
- g) an information theory based measure $I(I_u; I_w)$ representing mutual information between various phrases w and the user u;
- h) an information theory based measure $I(I_t; I_u)$ representing mutual information between various topics t and the user u;
- i) an information theory based measure $I(I_s; I_u)$ representing mutual information between various web sites s and the user u;
- j) an information theory based measure $I(I_p; I_u)$ representing mutual information between various products p and the user u; and
- k) an information theory based measure $I(I_f; I_u)$ representing mutual information between various features f of each of the various products p and the user u.

40. The program storage device of claim 32 wherein the parameters of the learning machine define:

- a) a user product probability distribution $P(p|u)$ representing interests of the user u in various products p; and
- b) a user product feature probability distribution $P(f|u,p)$ representing interests of the user u in various features f of each of the various products p;

and wherein the method steps further comprise estimating a probability $P(uld, \text{product described}=p)$ that a document d that describes a product p is of interest to the user u, wherein the probability is estimated in part the user product probability distribution and the user product feature probability distribution.

41. The program storage device of claim 40 wherein the method steps further comprise recommending products to the user based on the probability $P(uld, \text{product described}=p)$.

42. The program storage device of claim 32 wherein the method steps further comprise estimating a posterior probability $P(uld,q)$ that the document d is of interest to the user u, given a query q submitted by the user.

37

43. The program storage device of claim 42 wherein estimating the posterior probability comprises estimating a probability $P(q|d,u)$ that the query q is expressed by the user u with an information need in the document d .

44. The program storage device of claim 32 wherein the method steps further comprise applying the identified properties of the document d to a learning machine having product parameters characterizing a product p to estimate a probability $P(p|d)$ that the document d refers to the product p .

45. The program storage device of claim 44 wherein the method steps further comprise updating the product parameters based on the identified properties of the document d and the estimated probability $P(p|d)$.

46. The program storage device of claim 44 wherein the method steps further comprise initializing the product parameters based on a set of documents associated with the product p .

47. The program storage device of claim 32 wherein the method steps further comprise clustering multiple users into clusters of similar users, wherein the clustering comprises calculating distances between User Models, and selecting similar users based on the calculated distances between User Models.

48. The program storage device of claim 32 wherein the method steps further comprise calculating relative entropy values between User Models of multiple users, and clustering together users based on the calculated relative entropy values.

49. The program storage device of claim 32 wherein the parameters defining the User Model comprise calculated distances between the User Model and User Models of users similar to the user.

50. The program storage device of claim 32 wherein the method steps further comprise selecting in a group of users an expert user in an area of expertise, wherein selecting the expert user comprises finding an expert User Model among User Models of the group of users, such that the expert User Model indicates a strong interest of the expert user in a document associated with the area of expertise.

51. The program storage device of claim 32 wherein the method steps further comprise parsing the document d for hyperlinks, and separately estimating for each of the hyperlinks a probability that the hyperlink is of interest to the user u .

52. The program storage device of claim 32 wherein the method steps further comprise sending to a third party web

38

server user interest information derived from the User Model, whereby the third party web server may customize its interaction with the user.

53. The program storage device of claim 32 wherein the monitored user interactions include a sequence of interaction times.

54. The program storage device of claim 32 wherein the method steps further comprise initializing the User Model using information selected from the group consisting of a set of documents provided by the user, a web browser history file associated with the user, a web browser bookmarks file associated with the user, ratings by the user of a set of documents, and previous product purchases made by the user.

55. The program storage device of claim 32 wherein the method steps further comprise modifying the User Model based on User Model modification requests provided by the user.

56. The program storage device of claim 32 wherein the method steps further comprise providing to the user a score for a document identified by the user, wherein the score is derived from the estimated probability.

57. The program storage device of claim 32 wherein the method steps further comprise providing to the user a 3D map of a hyper linked document collection, wherein the 3D map indicates a user interest in each document.

58. The program storage device of claim 32 wherein the method steps further comprise temporarily using a User Model that is built from a set of predetermined parameters of a profile selected by the user.

59. The program storage device of claim 32 wherein the method steps further comprise initializing the User Model by selecting a set of predetermined parameters of a prototype user selected by the user.

60. The program storage device of claim 59 wherein the method steps further comprise updating the predetermined parameters of the prototype user based on actions of users similar to the prototype user.

61. The program storage device of claim 32 wherein the method steps further comprise identifying a set of users interested in the document d .

62. The program storage device of claim 61 wherein the method steps further comprise calculating a range of interests in the document d for the identified set of users.

* * * * *

EXHIBIT 2



US007685276B2

(12) **United States Patent**
Konig et al.

(10) **Patent No.:** **US 7,685,276 B2**
(45) **Date of Patent:** **Mar. 23, 2010**

(54) **AUTOMATIC, PERSONALIZED ONLINE INFORMATION AND PRODUCT SERVICES**

(75) Inventors: **Yochai Konig**, San Francisco, CA (US);
Roy Twersky, San Francisco, CA (US);
Michael R. Berthold, Berkeley, CA (US)

(73) Assignee: **Personalized User Model**, New York, NY (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **12/008,148**

(22) Filed: **Jan. 8, 2008**

(65) **Prior Publication Data**

US 2008/0114756 A1 May 15, 2008

Related U.S. Application Data

(63) Continuation of application No. 11/316,785, filed on Dec. 22, 2005, now Pat. No. 7,320,031, which is a continuation of application No. 09/597,975, filed on Jun. 20, 2000, now Pat. No. 6,981,040.

(60) Provisional application No. 60/173,392, filed on Dec. 28, 1999.

(51) **Int. Cl.**
G06F 15/16 (2006.01)

(52) **U.S. Cl.** **709/224**; 709/223; 709/228;
715/736

(58) **Field of Classification Search** 709/201–203,
709/223–228; 707/1–3, 7–10; 715/736–743
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,704,017 A 12/1997 Heckerman et al.

(Continued)

OTHER PUBLICATIONS

Pretschner, Alexander, "Ontology Based Personalized Search", *Master's Thesis, Department of Electrical Engineering and Computer Science*, University of Kansas, (1998), 125 pgs.

(Continued)

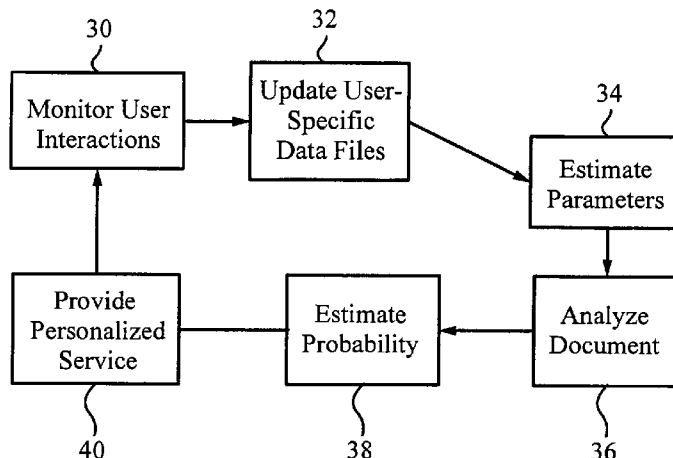
Primary Examiner—Bharat N Barot

(74) *Attorney, Agent, or Firm*—Schwegman, Lundberg & Woessner, P.A.

(57) **ABSTRACT**

A method for providing automatic, personalized information services to a computer user includes the following steps: transparently monitoring user interactions with data during normal use of the computer; updating user-specific data files including a set of user-related documents; estimating parameters of a learning machine that define a User Model specific to the user, using the user-specific data files; analyzing a document to identify its properties; estimating the probability that the user is interested in the document by applying the document properties to the parameters of the User Model; and providing personalized services based on the estimated probability. Personalized services include personalized searches that return only documents of interest to the user, personalized crawling for maintaining an index of documents of interest to the user; personalized navigation that recommends interesting documents that are hyperlinked to documents currently being viewed; and personalized news, in which a third party server customized its interaction with the user. The User Model includes continually-updated measures of user interest in words or phrases, web sites, topics, products, and product features. The measures are updated based on both positive examples, such as documents the user bookmarks, and negative examples, such as search results that the user does not follow. Users are clustered into groups of similar users by calculating the distance between User Models.

29 Claims, 19 Drawing Sheets



U.S. PATENT DOCUMENTS

5,754,939	A	5/1998	Herz et al.	
5,867,799	A	2/1999	Lang et al.	
5,918,014	A	6/1999	Robinson	
5,933,827	A	8/1999	Cole et al.	
5,964,839	A	10/1999	Johnson et al.	
5,983,214	A	11/1999	Lang	
5,991,735	A	11/1999	Gerace	
5,999,975	A	12/1999	Kittaka et al.	
6,006,218	A	12/1999	Breese et al.	
6,029,161	A	2/2000	Lang et al.	
6,041,311	A	3/2000	Chislenko et al.	
6,182,133	B1 *	1/2001	Horvitz	709/223
6,230,204	B1 *	5/2001	Fleming, III	709/224
6,564,170	B2	5/2003	Halabieh	
6,567,850	B1	5/2003	Freishtat et al.	
6,606,620	B1 *	8/2003	Sundaresan et al.	707/3
6,647,425	B1 *	11/2003	Chaddha	709/203
6,687,696	B2	2/2004	Hofmann et al.	
6,732,090	B2	5/2004	Shanahan et al.	
6,828,992	B1 *	12/2004	Freeman et al.	715/741
6,915,482	B2 *	7/2005	Jellum et al.	709/224
6,981,040	B1	12/2005	Konig et al.	
7,320,031	B2	1/2008	Konig	

OTHER PUBLICATIONS

"U.S. Appl. No. 09/597,975, Examiner Interview Summary mailed Dec. 11, 2003", 2 pgs.
 "U.S. Appl. No. 09/597,975, Final Office Action mailed Jun. 4, 2004", 32 pgs.
 "U.S. Appl. No. 09/597,975, Final Office Action mailed Dec. 3, 2003", 17 pgs.
 "U.S. Appl. No. 09/597,975, Non Final Office Action mailed Jan. 29, 2004", 25 pgs.

"U.S. Appl. No. 09/597,975, Non Final Office Action mailed Jun. 4, 2003", 13 pgs.
 "U.S. Appl. No. 09/597,975, Non Final Office Action mailed Jul. 8, 2005", 14 pgs.
 "U.S. Appl. No. 09/597,975, Non Final Office Action mailed Nov. 17, 2004", 30 pgs.
 "U.S. Appl. No. 09/597,975, Notice of Allowance mailed Sep. 22, 2005", 3 pgs.
 "U.S. Appl. No. 09/597,975, Response filed ar. 4, 2004 to Non Final Office Action mailed Jan. 29, 2004", 11 pgs.
 "U.S. Appl. No. 09/597,975, Response filed Aug. 4, 2005 to Non Final Office Action mailed Jul. 8, 2005", 7 pgs.
 "U.S. Appl. No. 09/597,975, Response filed Sep. 1, 2004 to Final Office Action mailed Jun. 4, 2004", 6 pgs.
 "U.S. Appl. No. 09/597,975, Response filed Sep. 4, 2003 to Non Final Office Action mailed 06-04-03", 10 pgs.
 "U.S. Appl. No. 09/597,975, Response filed Dec. 16, 2003 to Final Office Action mailed Dec. 3, 2003", 6 pgs.
 "U.S. Appl. No. 09/597,975, Response filed Dec. 22, 2004 to Non Final Office Action mailed Nov. 17, 2004", 24 pgs.
 "U.S. Appl. No. 11/316,785, Non Final Office Action mailed Aug. 22, 2006", 7 pgs.
 "U.S. Appl. No. 11/316,785, Notice of Allowance mailed Aug. 22, 2006", 3 pgs.
 "U.S. Appl. No. 11/316,785, Response filed Feb. 15, 2007 to Non Final Office Action mailed Aug. 22, 2006", 1 pg.
 Basset, et al., "A Study of Generalization Techniques in Evolutionary Rule Learning", *Paper*, (2002), 90 pgs.
 Lecun, Yann, "Machine Learning and Pattern Recognition", *lecture presented at The Courant Institute, New York University*, (2004).
 Mobasher, B., "Automatic personalization based on web usage mining", [Online]. Retrieved from the Internet: <URL:maya.cs.depaul.edu/~mobasher/personalization>, 25 pgs.

* cited by examiner

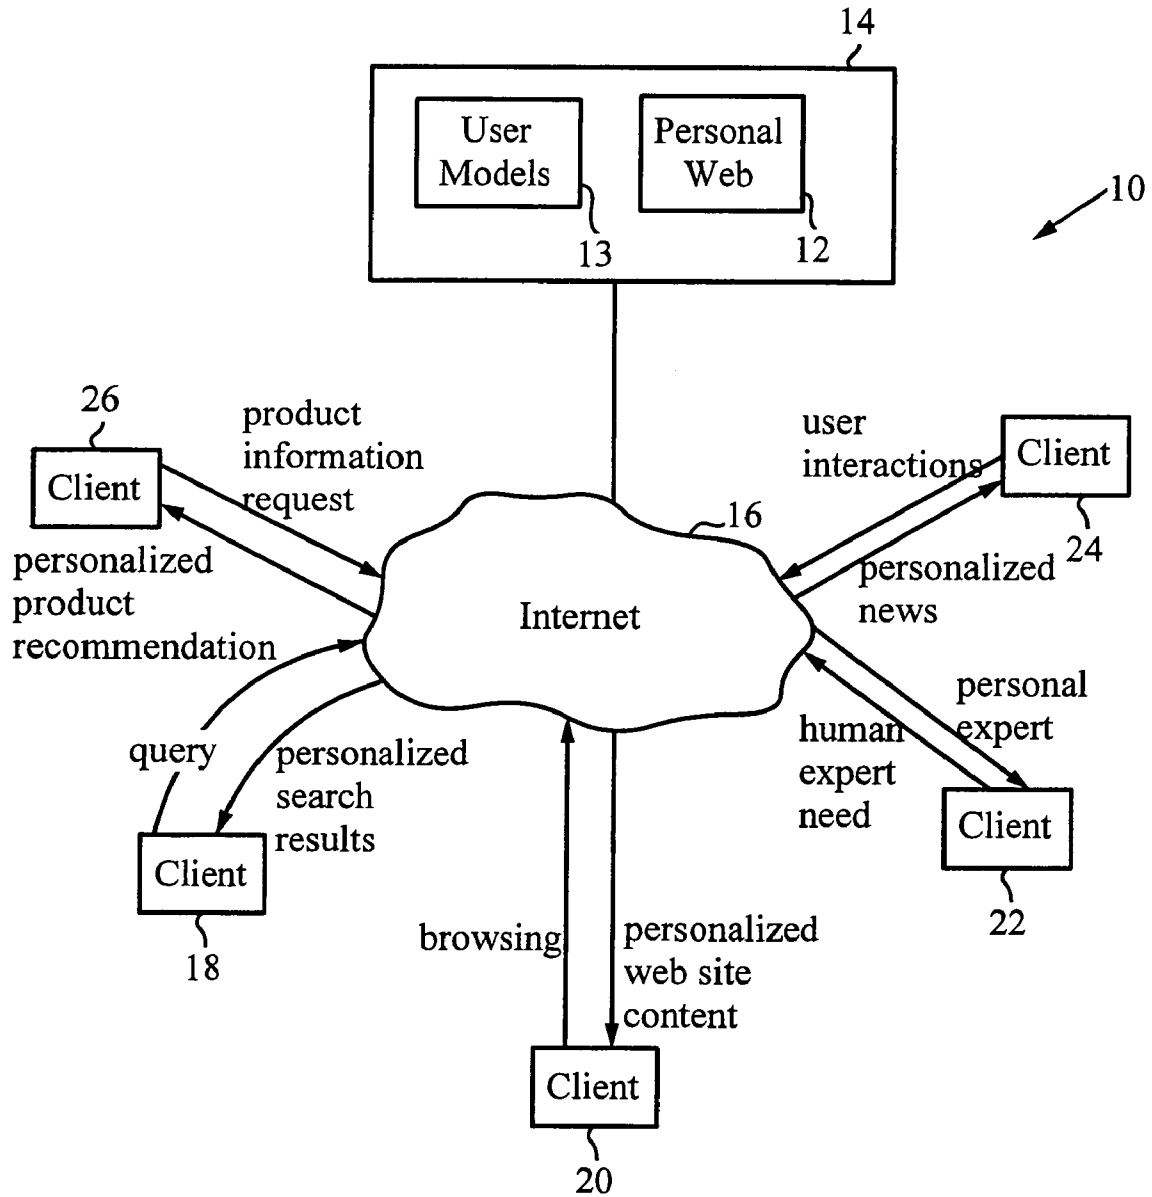


Fig. 1

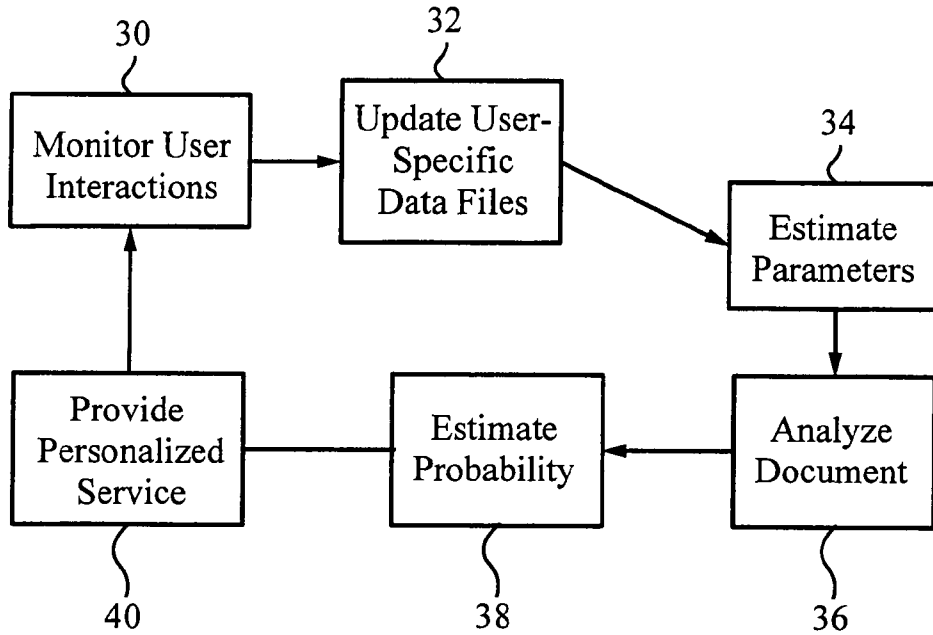


Fig. 2

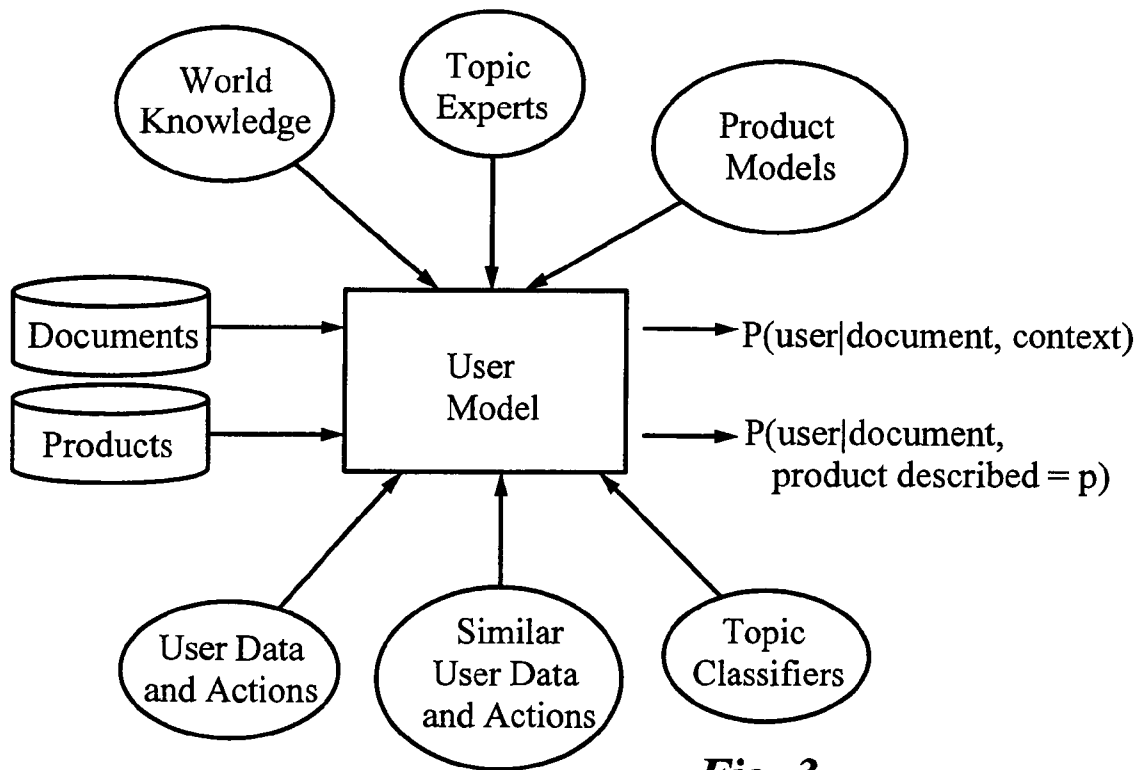


Fig. 3

Informative Word/Phrase List

Word ID	Word Grade	Last Access Time	Number of Accesses
Vegan	0.86	3/6/2000 12:22:41	173
Parasail	0.72	4/15/2000 18:51:27	220

Fig. 4A

Web Site Distribution

Site ID	Site Probability	Last Access Time	Number of Accesses
herring.com	0.61	5/1/2000 19:15:21	152
Java.com	0.43	4/24/2000 3:16:18	460

Fig. 4B

User Topic Distribution

Topic ID	Topic Parent	Topic Probability	Last Access Time	Number of Accesses
Computers	Industries	0.6	12/2/1999 1:21:22	74
Publishing	Industries	0.31	1/2/2000 6:25:31	62

Fig. 4C

User Product Distribution

Product ID	Product Parent	Product Probability	Last Purchase Time	Number of Purchases	Last Access Time	Number of Accesses
3Com Palm 3E	Without Keyboards	0.73	12/16/1999 17:21:21	1	5/2/2000 16:01:21	78
Without Keyboards	Handhelds/PDAs	0.81	12/16/1999 17:21:21	1	3/15/2000 17:21:21	98

Fig. 4D

User Product Feature Distribution

Product ID	Feature ID	Value	Value Probability
Webcams	Interface	PC Card	0.7
Webcams	Interface	Serial	0.2

Fig. 4E

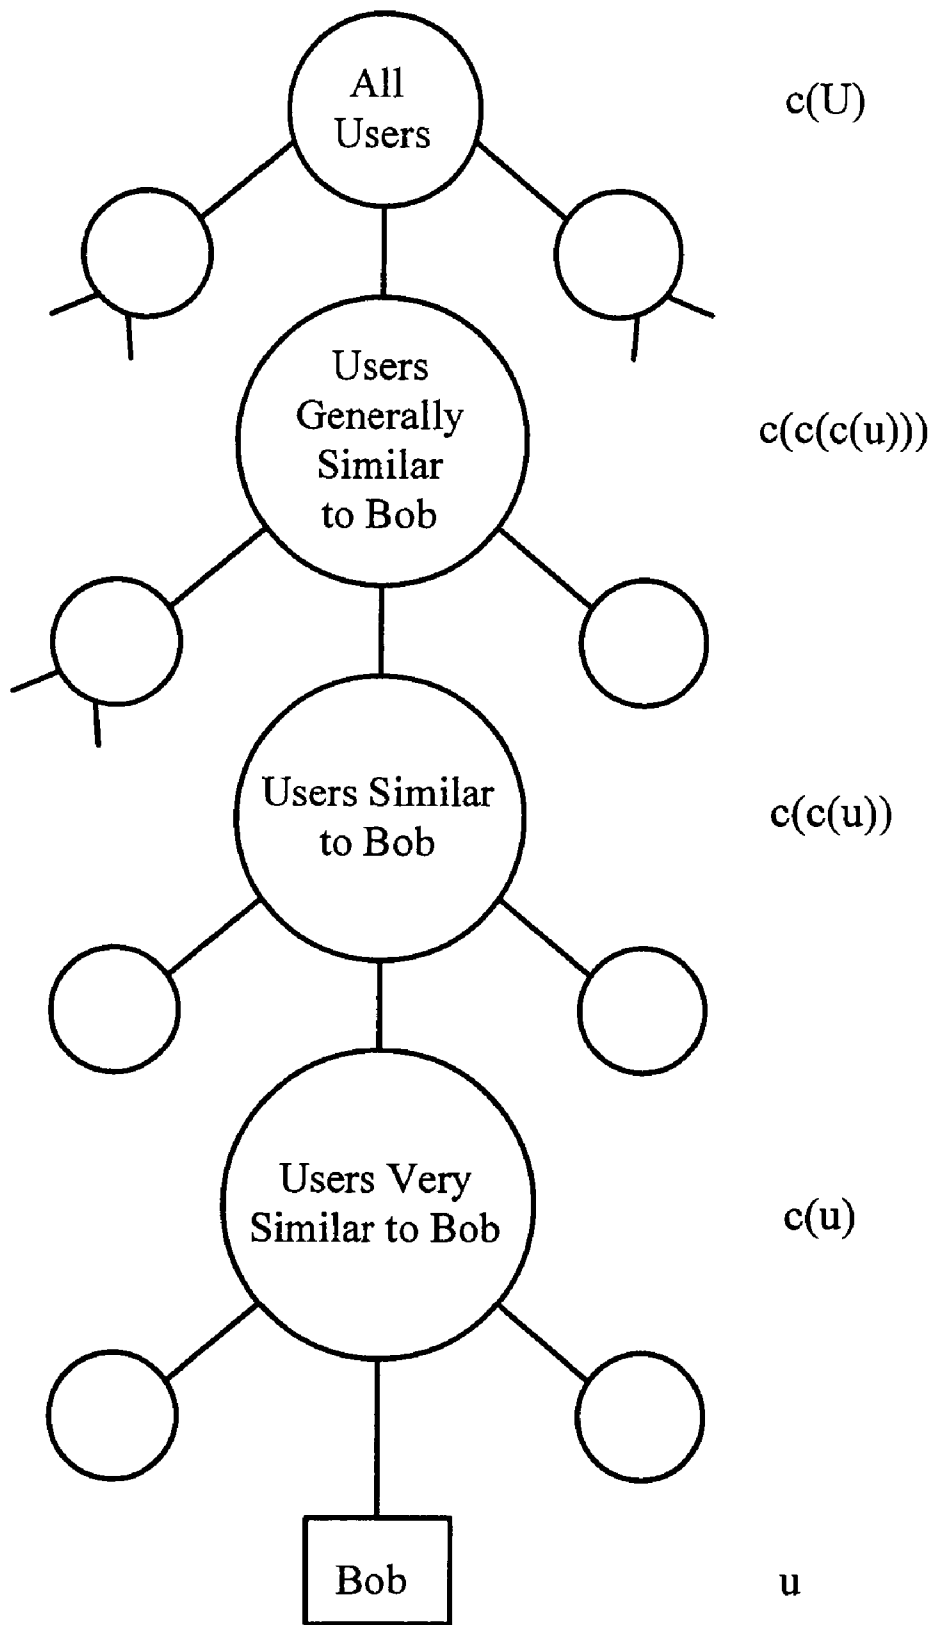


Fig. 5A

User Cluster Tree

Cluster ID	Cluster Parent ID
C123	C3345

Fig. 5B




Fig. 6A

User Fuzzy Cluster Tree

Cluster ID	Cluster Parent ID	Cluster Probability
Bob	C1	0.3
Bob	C2	0.2
Bob	C3	0.1
Bob	C4	0.4
C1	C11	0.2

Fig. 6B

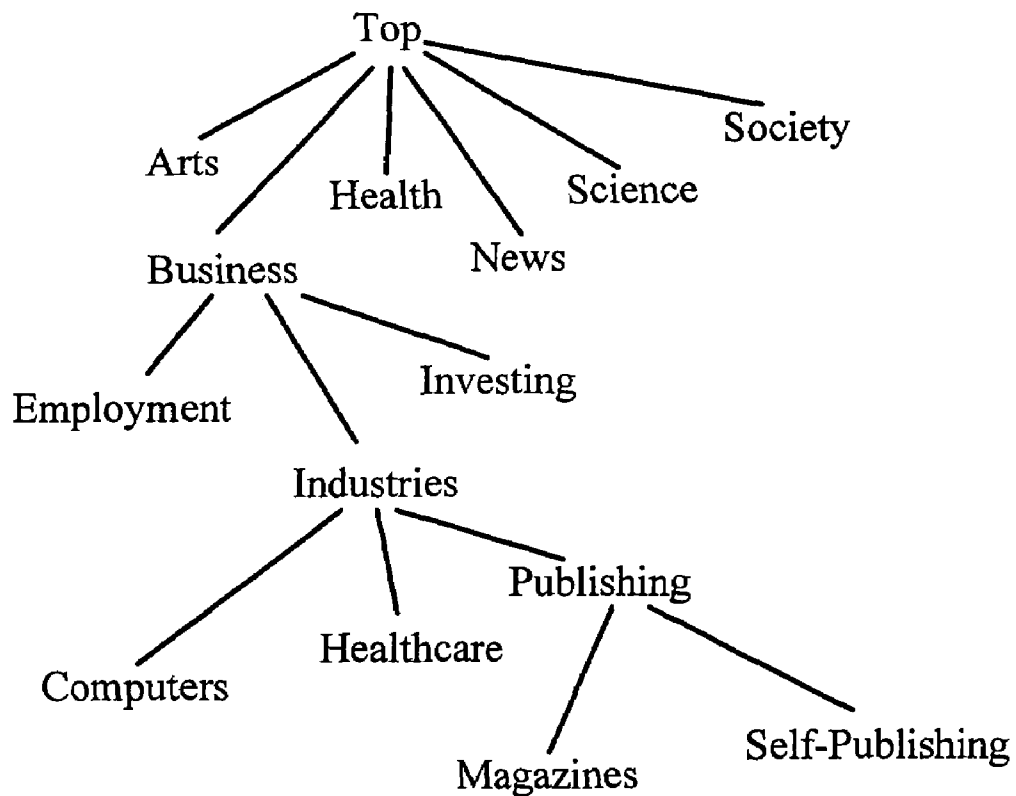


Fig. 7

Topic Tree

Topic ID	Depth Level	Topic Parent ID	Number of Children	Children
Investing	2	Business	3	International, Resources, Socially Responsible
Employment	2	Business	4	Careers, Recruiters, Resumes, Seasonal

Fig. 8

Topic Experts

Topic ID	Topic Parent ID	Cluster 1	Cluster 2
Investing	Business	C112	C113 C114
Employment	Business	C241	C212 C159

Fig. 9




Fig. 10

Product Tree

Product ID	Depth Level	Product Parent ID	Number of Children	Children
Cameras	3	Consumer Electronics	2	Digital Cameras, Webcams
Consumer Electronics	2	Top	3	CD Players, Cameras, Personal Minidisks

Fig. 11

Product Feature List

Product ID	Feature	Value
Sony CDP-CX350	Brand	Sony
Sony CDP-CX350	CD Capacity	50 Discs or Greater
Sony CDP-CX350	Digital Output	Optical

Fig. 12A

Product Feature Value List

Feature	Value
Digital Output	Coaxial and Optical
Digital Output	Coaxial
Digital Output	Optical
Digital Output	No

Fig. 12B




Fig. 13

User Recently Accessed Butter

Document ID	Access Time	Interaction Type	Context	Degree of Interest
www.herring.com/insider	5/12/2000 14:37:21	Navigation	bookmark access	positive, followed 3 links 12 minutes
www.m-w.com	5/12/2000 15:08:21	Search	query "dictionary"	positive, followed 5 links bookmarked, 21 minutes

Fig. 14

User Site Candidate Table

Site Name	Number of Access	Last Access Time
www.herring.com	157	5/12/2000 14:37:21
www.m-w.com	162	5/12/2000 15:08:21

Fig. 15A

User Word Candidate Table

Word ID	Word Spelling	Word Spelling	Word Grade	Last Access Time
Cytochrome	Cytochrome	Cytocrome	0.67	4/16/2000 7:10:01
Hyperbilirubinemia	Hyperbilirubinemia	Hyperbilirubenema	0.58	4/27/2000 12:18:42

Fig. 15B

User Recently Purchased Products

Product ID	Parent Node	Purchase Time	Purchase Source
Panasonic SL-502	Discmans	5/1/2000 16:01:04	ebyweb.com
Hitachi VM6500A	Camcorders	5/3/2000 18:19:21	supremevideo.com

Fig. 16

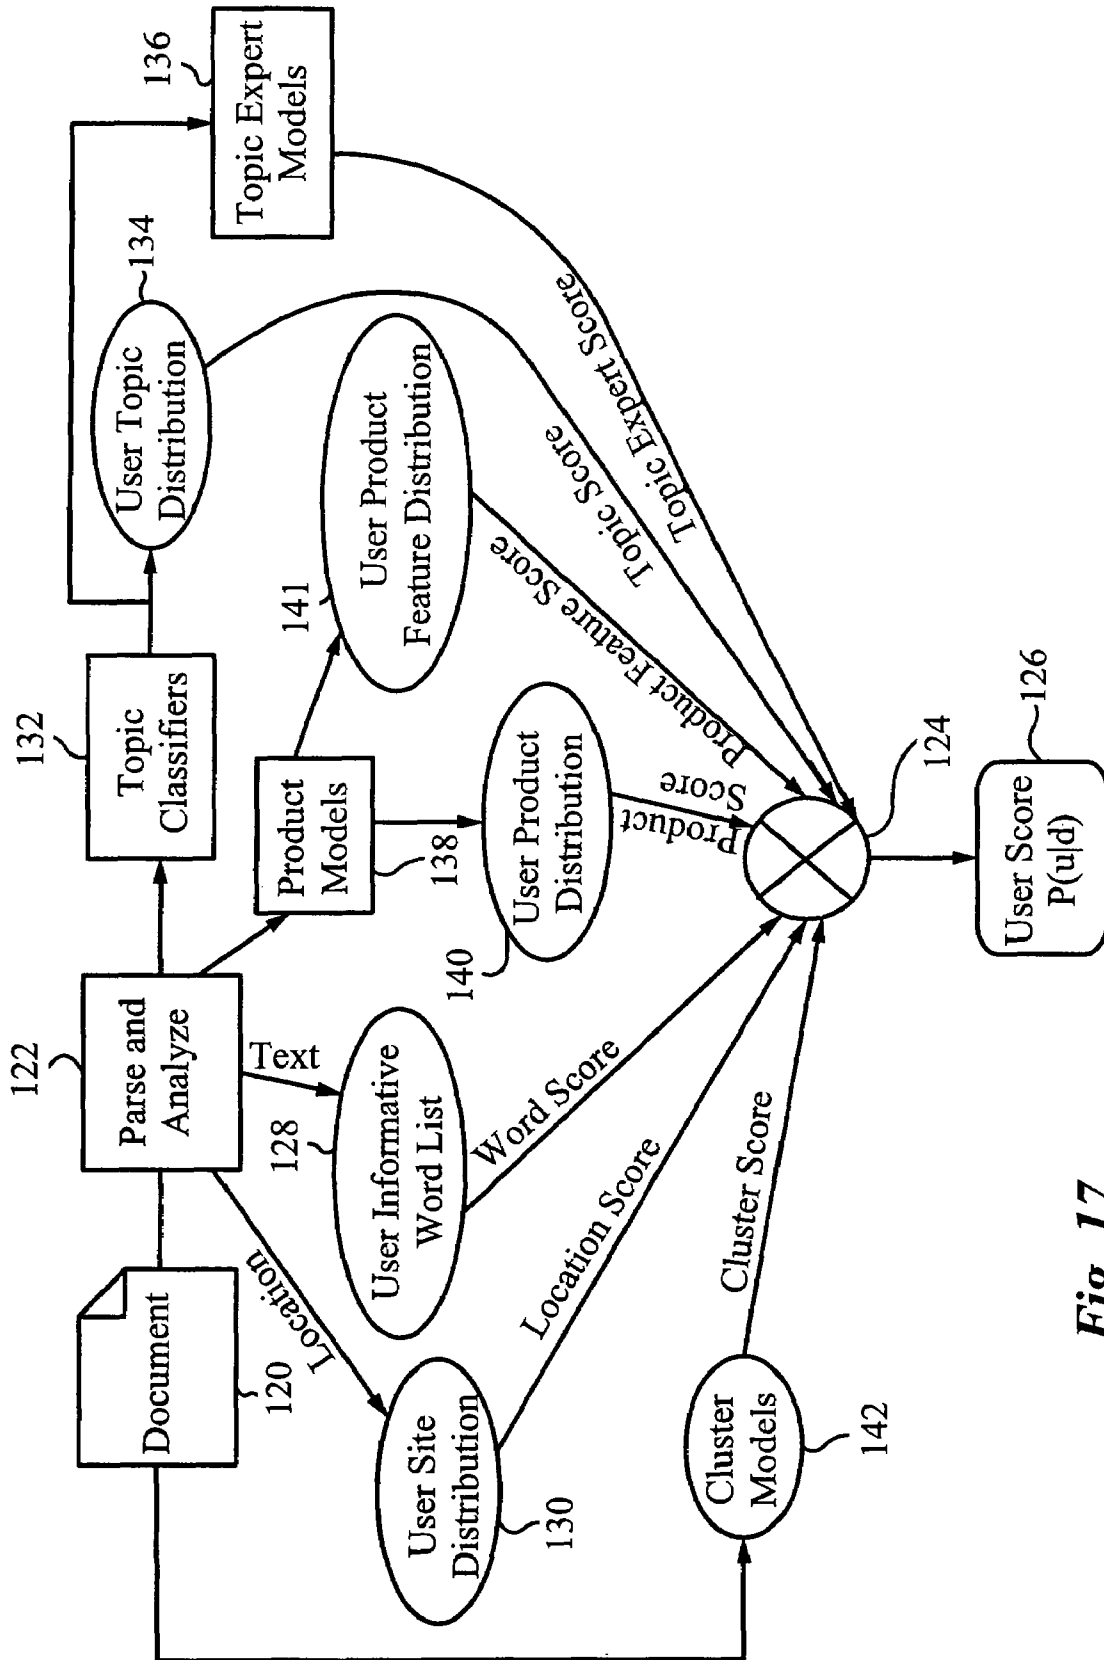


Fig. 17

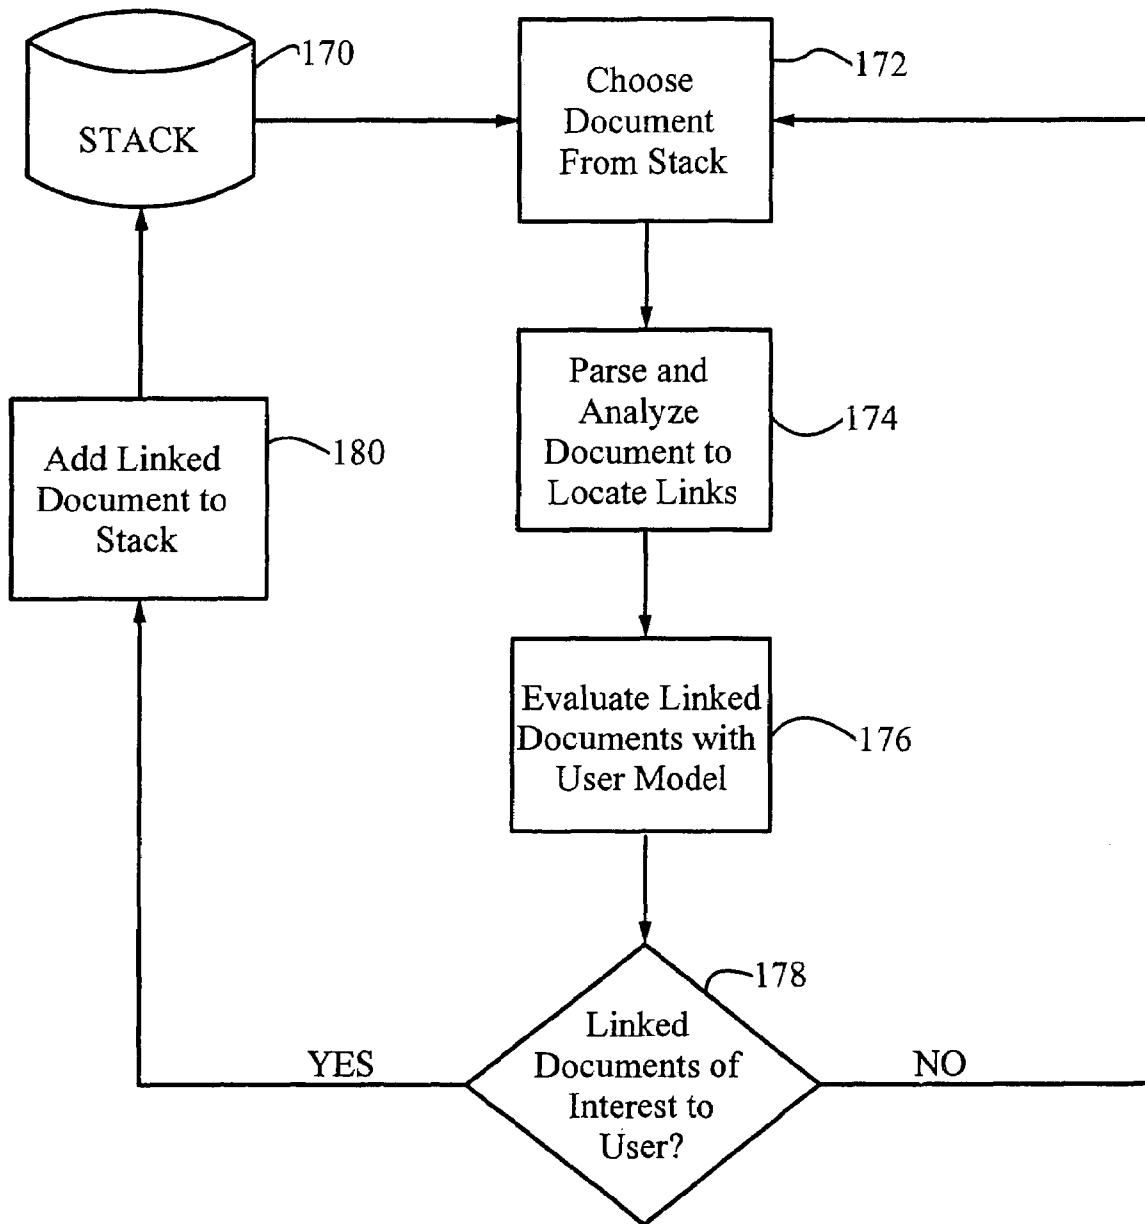


Fig. 18

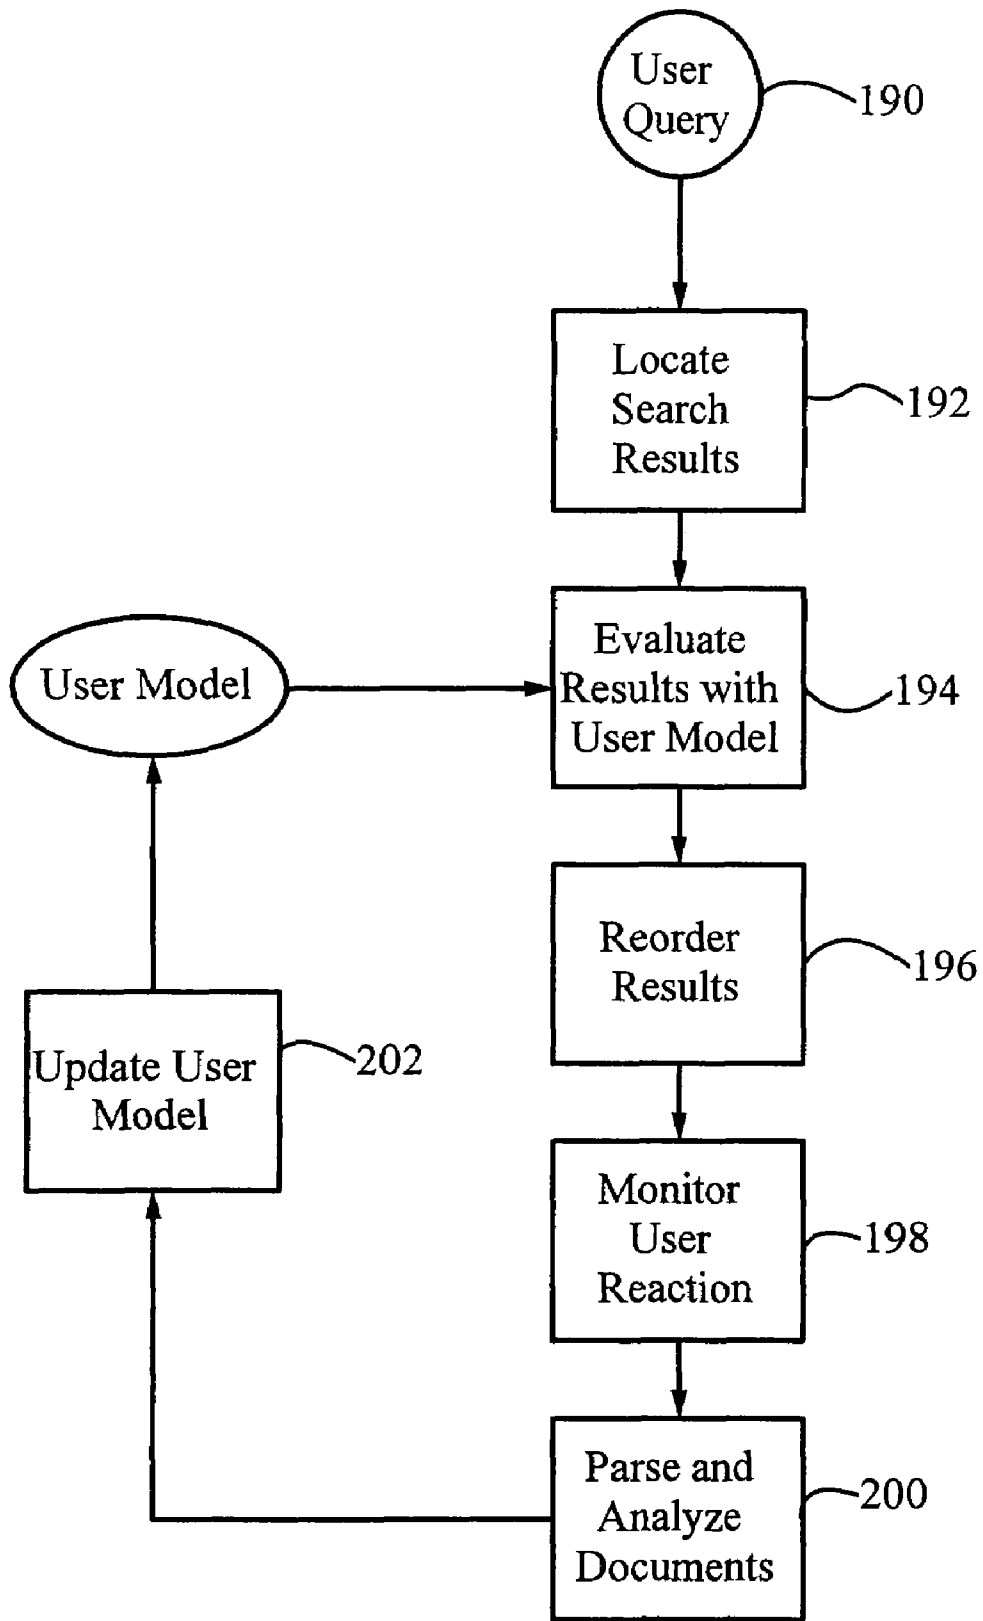


Fig. 19

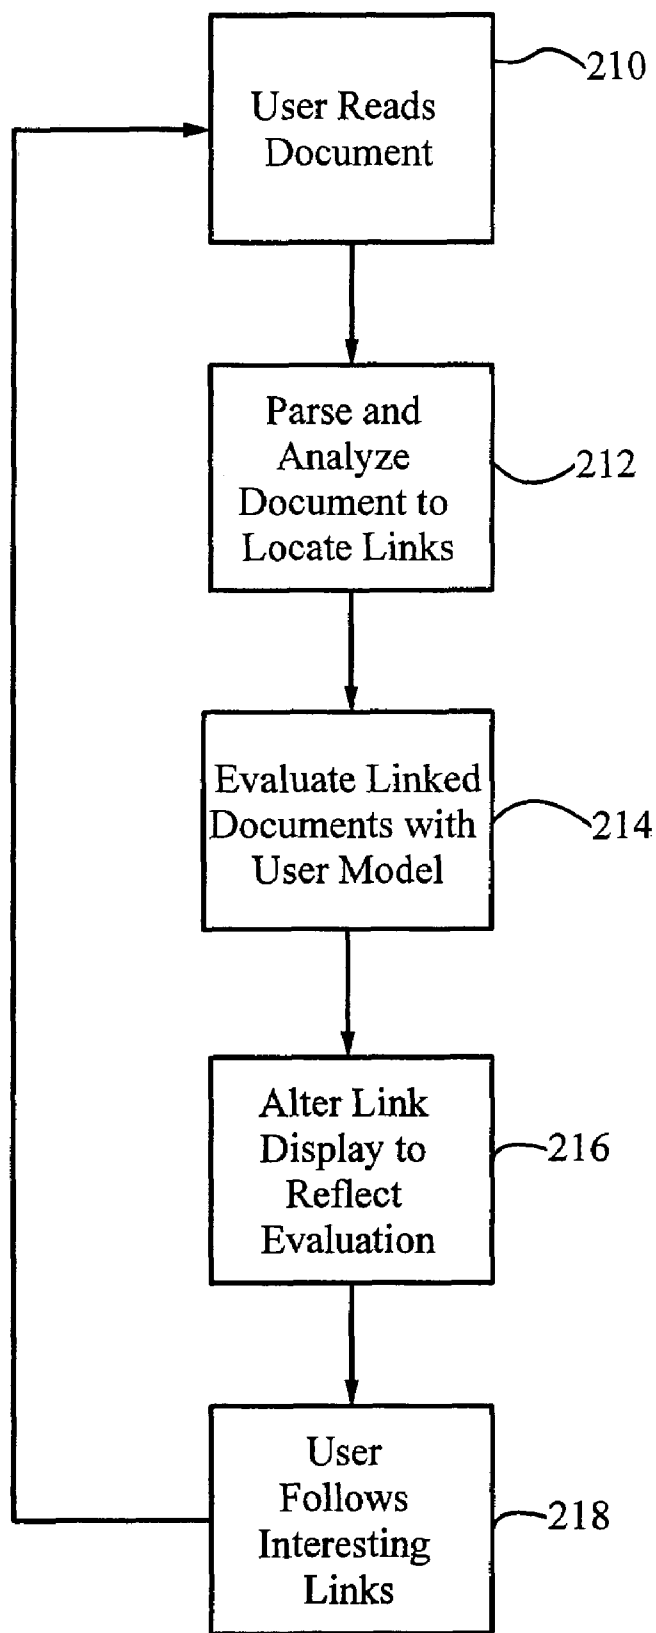


Fig. 20

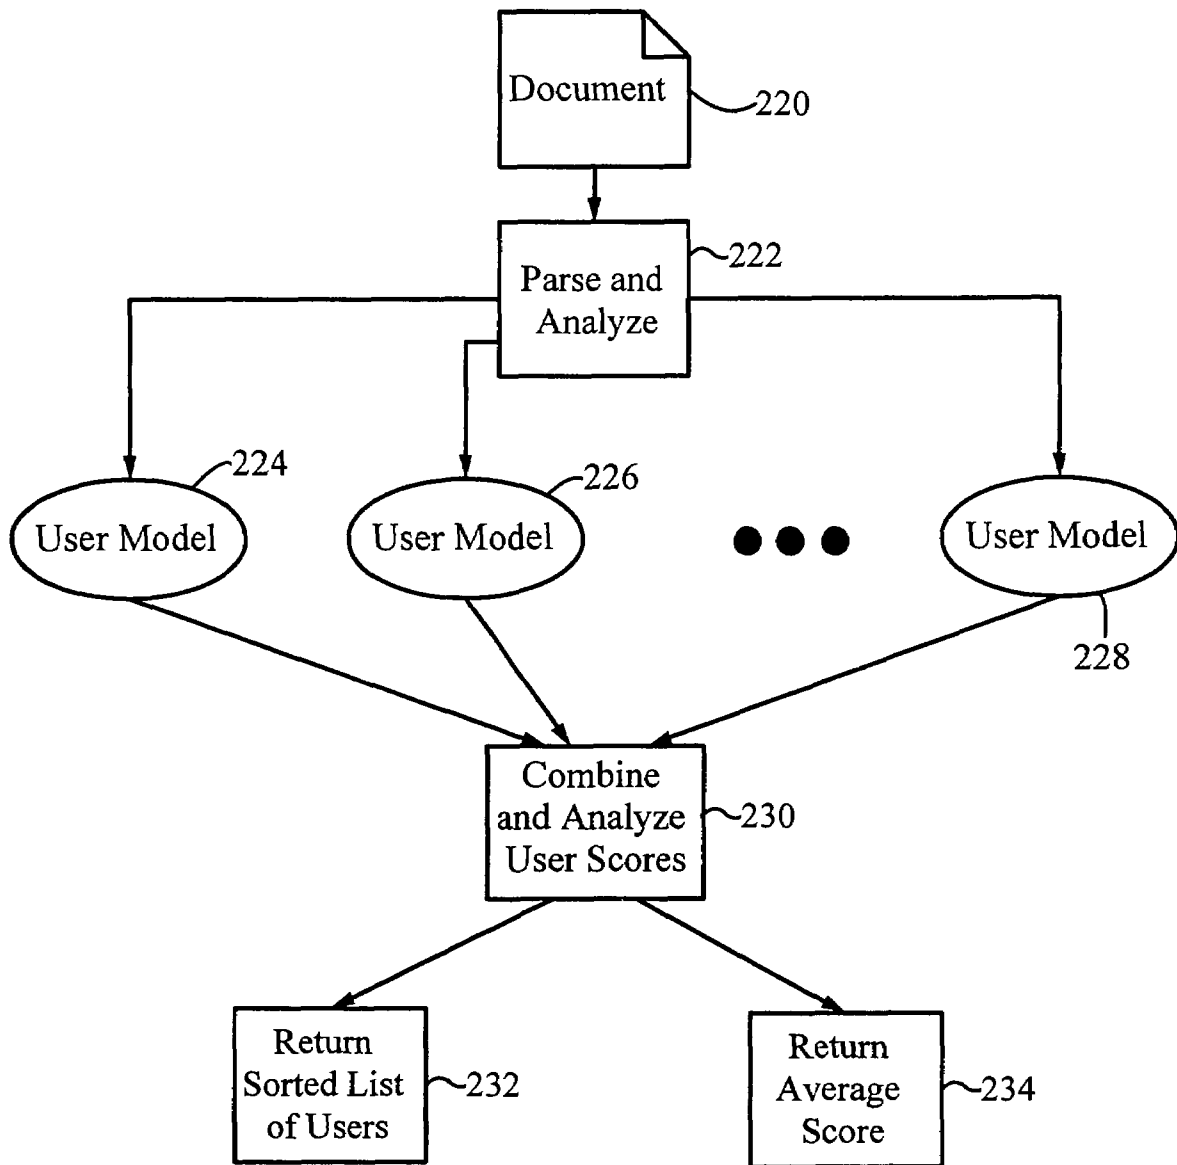


Fig. 21




Fig. 22

AUTOMATIC, PERSONALIZED ONLINE INFORMATION AND PRODUCT SERVICES

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation application of U.S. Non-Provisional application Ser. No. 11/316,785 filed Dec. 22, 2005 now U.S. Pat. No. 7,320,031, which is a continuation application of U.S. Non-Provisional application Ser. No. 09/597,975 filed Jun. 20, 2000 now U.S. Pat. No. 6,981,040. Accordingly, this application claims the benefit of U.S. Non-Provisional application Ser. No. 09/597,975 filed Jun. 20, 2000, which claims the benefit of U.S. Provisional Application No. 60/173,392 filed Dec. 28, 1999, which are all herein incorporated by reference.

FIELD OF THE INVENTION

This invention relates generally to methods for personalizing a user's interaction with information in a computer network. More particularly, it relates to methods for predicting user interest in documents and products using a learning machine that is continually updated based on actions of the user and similar users.

BACKGROUND ART

The amount of static and dynamic information available today on the Internet is staggering, and continues to grow exponentially. Users searching for information, news, or products and services are quickly overwhelmed by the volume of information, much of it useless and uninformative. A variety of techniques have been developed to organize, filter, and search for information of interest to a particular user. Broadly, these methods can be divided into information filtering techniques and collaborative filtering techniques.

Information filtering techniques focus on the analysis of item content and the development of a personal user interest profile. In the simplest case, a user is characterized by a set of documents, actions regarding previous documents, and user-defined parameters, and new documents are characterized and compared with the user profile. For example, U.S. Pat. No. 5,933,827, issued to Cole et al., discloses a system for identifying new web pages of interest to a user. The user is characterized simply by a set of categories, and new documents are categorized and compared with the user's profile. U.S. Pat. No. 5,999,975, issued to Kittaka et al., describes an online information providing scheme that characterizes users and documents by a set of attributes, which are compared and updated base on user selection of particular documents. U.S. Pat. No. 6,006,218, issued to Breese et al., discloses a method for retrieving information based on a user's knowledge, in which the probability that a user already knows of a document is calculated based on user-selected parameters or popularity of the document. U.S. Pat. No. 5,754,939, issued to Herz et al., discloses a method for identifying objects of interest to a user based on stored user profiles and target object profiles. Other techniques rate documents using the TFIDF (term frequency, inverse document frequency) measure. The user is represented as a vector of the most informative words in a set of user-associated documents. New documents are parsed to obtain a list of the most informative words, and this list is compared to the user's vector to determine the user's interest in the new document.

Existing information filtering techniques suffer from a number of drawbacks. Information retrieval is typically a two

step process, collection followed by filtering; information filtering techniques personalize only the second part of the process. They assume that each user has a personal filter, and that every network document is presented to this filter. This assumption is simply impractical given the current size and growth of the Internet; the number of web documents is expected to reach several billion in the next few years. Furthermore, the dynamic nature of the documents, e.g., news sites that are continually updated, makes collection of documents to be filtered later a challenging task for any system. User representations are also relatively limited, for example, including only a list of informative words or products or user-chosen parameters, and use only a single mode of interaction to make decisions about different types of documents and interaction modes. In addition, information filtering techniques typically allow for extremely primitive updating of a user profile, if any at all, based on user feedback to recommended documents. As a user's interests change rapidly, most systems are incapable of providing sufficient personalization of a user's experience.

Collaborative filtering methods, in contrast, build databases of user opinions of available items, and then predict a user opinion based on the judgments of similar users. Predictions typically require offline data mining of very large databases to recover association rules and patterns; a significant amount of academic and industrial research is focussed on developing more efficient and accurate data mining techniques. The earliest collaborative filtering systems required explicit ratings by the users, but existing systems are implemented without the user's knowledge by observing user actions. Ratings are inferred from, for example, the amount of time a user spends reading a document or whether a user purchases a particular product. For example, an automatic personalization method is disclosed in B. Mobasher et al., "Automatic Personalization Through Web Usage Mining," Technical Report TR99-010, Department of Computer Science, Depaul University, 1999. Log files of documents requested by users are analyzed to determine usage patterns, and online recommendations of pages to view are supplied to users based on the derived patterns and other pages viewed during the current session.

Recently, a significant number of web sites have begun implementing collaborative filtering techniques, primarily for increasing the number and size of customer purchases. For example, Amazon.com™ has a "Customers Who Bought" feature, which recommends books frequently purchased by customers who also purchased a selected book, or authors whose work is frequently purchased by customers who purchased works of a selected author. This feature uses a simple "shopping basket analysis"; items are considered to be related only if they appear together in a virtual shopping basket. Net Perceptions, an offshoot of the GroupLens project at the University of Minnesota, is a company that provides collaborative filtering to a growing number of web sites based on data mining of server logs and customer transactions, according to predefined customer and product clusters.

Numerous patents disclose improved collaborative filtering systems. A method for item recommendation based on automated collaborative filtering is disclosed in U.S. Pat. No. 6,041,311, issued to Chislenko et al. Similarity factors are maintained for users and for items, allowing predictions based on opinions of other users. In an extension of standard collaborative filtering, item similarity factors allow predictions to be made for a particular item that has not yet been rated, but that is similar to an item that has been rated. A method for determining the best advertisements to show to users is disclosed in U.S. Pat. No. 5,918,014, issued to Rob-

inson. A user is shown a particular advertisement based on the response of a community of similar users to the particular advertisement. New ads are displayed randomly, and the community interest is recorded if enough users click on the ads. A collaborative filtering system using a belief network is disclosed in U.S. Pat. No. 5,704,317, issued to Heckerman et al., and allows automatic clustering and use of non-numeric attribute values of items. A multi-level mindpool system for collaborative filtering is disclosed in U.S. Pat. No. 6,029,161, issued to Lang et al. Hierarchies of users are generated containing clusters of users with similar properties.

Collaborative filtering methods also suffer from a number of drawbacks, chief of which is their inability to rate content of an item or incorporate user context. They are based only on user opinions; thus an item that has never been rated cannot be recommended or evaluated. Similarly, obscure items, which are rated by only a few users, are unlikely to be recommended. Furthermore, they require storage of a profile for every item, which is unfeasible when the items are web pages. New items cannot be automatically added into the database. Changing patterns and association rules are not incorporated in real time, since the data mining is performed offline. In addition, user clusters are also static and cannot easily be updated dynamically.

Combinations of information filtering and collaborative filtering techniques have the potential to supply the advantages provided by both methods. For example, U.S. Pat. No. 5,867,799, issued to Lang et al., discloses an information filtering method that incorporates both content-based filtering and collaborative filtering. However, as with content-based methods, the method requires every document to be filtered as it arrives from the network, and also requires storage of a profile of each document. Both of these requirements are unfeasible for realistically large numbers of documents. An extension of this method, described in U.S. Pat. No. 5,983,214, also to Lang et al., observes the actions of users on content profiles representing information entities. Incorporating collaborative information requires that other users have evaluated the exact content profile for which a rating is needed.

In summary, none of the existing prior art methods maintain an adaptive content-based model of a user that changes based on user behavior, allow for real-time updating of the model, operate during the collection stage of information retrieval, can make recommendations for items or documents that have never been evaluated, or model a user based on different modes of interaction.

OBJECTS AND ADVANTAGES

Accordingly, it is a primary object of the present invention to provide a method of personalizing user interaction with network documents that maintains an adaptive content-based profile of the user.

It is another object of the invention to incorporate into the profile user behavior during different modes of interaction with information, thus allowing for cross-fertilization. Learning about the user interests in one mode benefits all other modes.

It is a further object of the invention to provide a method that jointly models the user's information needs and product needs to provide stronger performance in both modes.

It is an additional object of the invention to provide a method that personalizes both the collection and filtering stages of information retrieval to manage efficiently the enormous number of existing web documents.

It is another object of the invention to provide a method for predicting user interest in an item that incorporates the opinions of similar users without requiring storage and maintenance of an item profile.

It is a further object of the invention to provide an information personalization method that models the user as a function independent of any specific representation or data structure, and represents the user interest in a document or product independently of any specific user information need. This approach enables the addition of new knowledge sources into the user model.

It is an additional object of the present invention to provide a method based on Bayesian statistics that updates the user profile based on both negative and positive examples.

It is a further object of the invention to model products by analyzing all relevant knowledge sources, such as press releases, reviews, and articles, so that a product can be recommended even if it has never been purchased or evaluated previously.

SUMMARY

These objects and advantages are attained by a computer-implemented method for providing automatic, personalized information services to a user. User interactions with a computer are transparently monitored while the user is engaged in normal use of the computer, and monitored interactions are used to update user-specific data files that include a set of documents associated with the user. Parameters of a learning machine, which define a User Model specific to the user, are estimated from the user-specific data files. Documents that are of interest and documents that are not of interest to the user are treated distinctly in estimating the parameters. The parameters are used to estimate a probability $P(u|d)$ that a document is of interest to the user, and the estimated probability is then used to provide personalized information services to the user.

The probability is estimated by analyzing properties of the document and applying them to the learning machine. Documents of multiple distinct media types of analyzed, and identified properties include: the probability that the document is of interest to users who are interested in particular topics, a topic classifier probability distribution, a product model probability distribution, product feature values extracted from the document, the document author, the document age, a list of documents linked to the document, the document language, number of users who have accessed the document, number of users who have saved the document in a favorite document list, and a list of users previously interested in the document. All properties are independent of the particular user. The product model probability distribution, which indicates the probability that the document refers to particular products, is obtained by applying the document properties to a product model, a learning machine with product parameters characterizing particular products. These product parameters are themselves updated based on the document properties and on the product model probability distribution. Product parameters are initialized from a set of documents associated with each product.

User interactions are monitored during multiple distinct modes of user interaction with network data, including a network searching mode, network navigation mode, network browsing mode, email reading mode, email writing mode, document writing mode, viewing "pushed" information mode, finding expert advice mode, and product purchasing mode. Based on the monitored interactions, parameters of the learning machine are updated. Learning machine parameters

define various user-dependent functions of the User Model, including a user topic probability distribution representing interests of the user in various topics, a user product probability distribution representing interests of the user in various products, a user product feature probability distribution representing interests of the user in various features of each of the various products, a web site probability distribution representing interests of the user in various web sites, a cluster probability distribution representing similarity of the user to users in various clusters, and a phrase model probability distribution representing interests of the user in various phrases. Some of the user-dependent functions can be represented as information theory based measures representing mutual information between the user and either phrases, topics, products, features, or web sites. The product and feature distributions can also be used to recommend products to the user.

The User Model is initialized from documents provided by the user, a web browser history file, a web browser bookmarks file, ratings by the user of a set of documents, or previous product purchases made by the user. Alternatively, the User Model may be initialized by selecting a set of predetermined parameters of a prototype user selected by the user. Parameters of the prototype user are updated based on actions of users similar to the prototype user. The User Model can be modified based on User Model modification requests provided by the user. In addition, the user can temporarily use a User Model that is built from a set of predetermined parameters of a profile selected by the user.

Distances between users are calculated to determine similar users, who are clustered into clusters of similar users. Parameters defining the User Model may include the calculated distances between the User Model and User Models of users within the user's cluster. Users may also be clustered based on calculated relative entropy values between User Models of multiple users.

A number of other probabilities can be calculated, such as a posterior probability $P(u|d,q)$ that the document is of interest to the user, given a search query submitted by the user. Estimating the posterior probability includes estimating a probability that the query is expressed by the user with an information need contained in the document. In addition, the probability $P(u|d,con)$ that the document is of interest to the user during a current interaction session can be calculated. To do so, $P(u,con|d)/P(con|d)$ is calculated, where con represents a sequence of interactions during the current interaction session or media content currently marked by the user. A posterior probability $P(u|d,q,con)$ that the document is of interest to the user, given a search query submitted during a current interaction session, can also be calculated.

A variety of personalized information services are provided using the estimated probabilities. In one application, network documents are crawled and parsed for links, and probable interest of the user in the links is calculated using the learning machine. Links likely to be of interest to the user are followed. In another application, the user identifies a document, and a score derived from the estimated probability is provided to the user. In an additional application, the user is provided with a three-dimensional map indicating user interest in each document of a hyperlinked document collection. In a further application, an expert user is selected from a group of users. The expert user has an expert User Model that indicates a strong interest in a document associated with a particular area of expertise. Another application includes parsing a viewed document for hyperlinks and separately estimating for each hyperlink a probability that the linked document is of interest to the user. In a further application,

user interest information derived from the User Model is sent to a third party web server that then customizes its interaction with the user. Finally, a set of users interested in a document is identified, and a range of interests for the identified users is calculated.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 is a schematic diagram of a computer system in which the present invention is implemented.

FIG. 2 is a block diagram of a method of the present invention for providing personalized product and information services to a user.

FIG. 3 is a schematic diagram of knowledge sources used as inputs to the User Model and resulting outputs.

FIGS. 4A-4E illustrate tables that store different components and parameters of the User Model.

FIG. 5A illustrates a cluster tree containing clusters of users similar to a particular user.

FIG. 5B is a table that stores parameters of a user cluster tree.

FIG. 6A illustrates a preferred cluster tree for implementing fuzzy or probabilistic clustering.

FIG. 6B is a table that stores parameters of a user fuzzy cluster tree.

FIG. 7 illustrates a portion of a topic tree.

FIG. 8 is a table that stores nodes of the topic tree of FIG. 7.

FIG. 9 is a table that stores the names of clusters having the most interest in nodes of the topic tree of FIG. 7, used to implement the topic experts model.

FIG. 10 illustrates a portion of a product tree.

FIG. 11 is a table that stores nodes of the product tree of FIG. 10.

FIG. 12A is a table that stores feature values of products of the product tree of FIG. 10.

FIG. 12B is a table that stores potential values of product features associated with intermediate nodes of the product tree of FIG. 10.

FIG. 13 is a schematic diagram of the method of initializing the User Model.

FIG. 14 illustrates the user recently accessed buffer, which records all user interactions with documents.

FIG. 15A is a table for storing sites that are candidates to include in the user site distribution.

FIG. 15B is a table for storing words that are candidates to include in the user word distribution.

FIG. 16 is a table that records all products the user has purchased.

FIG. 17 is a schematic diagram of the method of applying the User Model to new documents to estimate the probability of user interest in the document.

FIG. 18 is a block diagram of the personal crawler application of the present invention.

FIG. 19 is a block diagram of the personal search application of the present invention.

FIG. 20 is a block diagram of the personal navigation application of the present invention.

FIG. 21 is a block diagram of the document barometer application of the present invention.

FIG. 22 is a schematic diagram of the three-dimensional map application of the present invention.

DETAILED DESCRIPTION

Although the following detailed description contains many specifics for the purposes of illustration, anyone of ordinary

skill in the art will appreciate that many variations and alterations to the following details are within the scope of the invention. Accordingly, the following preferred embodiment of the invention is set forth without any loss of generality to, and without imposing limitations upon, the claimed invention.

The present invention, referred to as Personal Web, provides automatic, personalized information and product services to a computer network user. In particular, Personal Web is a user-controlled, web-centric service that creates for each user a personalized perspective and the ability to find and connect with information on the Internet, in computer networks, and from human experts that best matches his or her interests and needs. A computer system **10** implementing Personal Web **12** is illustrated schematically in FIG. **1**. Personal Web **12** is stored on a central computer or server **14** on a computer network, in this case the Internet **16**, and interacts with client machines **18**, **20**, **22**, **24**, **26** via client-side software. Personal Web **12** may also be stored on more than one central computers or servers that interact over the network. The client-side software may be part of a web browser, such as Netscape Navigator or Microsoft Internet Explorer, configured to interact with Personal Web **12**, or it may be distinct from but interacting with a client browser. Five client machines are illustrated for simplicity, but Personal Web **12** is intended to provide personalized web services for a large number of clients simultaneously.

For all of the typical interactions that a user has with a computer network, such as the world wide web, Personal Web **12** provides a personalized version. Personal Web **12** stores for each user a User Model **13** that is continuously and transparently updated based on the user's interaction with the network, and which allows for personalization of all interaction modes. The User Model represents the user's information and product interests; all information that is presented to the user has been evaluated by the User Model to be of interest to the user. The User Model allows for cross fertilization; that is, information that is learned in one mode of interaction is used to improve performance in all modes of interaction. The User Model is described in detail below.

Five examples of personalized interaction modes provided by the present invention are illustrated in FIG. **1**. However, it is to be understood that the present invention provides for personalization of all modes, and that the following examples in no way limit the scope of the present invention. Personal Web is active during all stages of information processing, including collection, retrieval, filtering, routing, and query answering.

Client **18** performs a search using Personal Web **12** by submitting a query and receiving personalized search results. The personal search feature collects, indexes, and filters documents, and responds to the user query, all based on the user profile stored in the User Model **13**. For example, the same query (e.g., "football game this weekend" or "opera") submitted by a teenager in London and an adult venture capitalist in Menlo Park returns different results based on the personality, interests, and demographics of each user. By personalizing the collection phase, the present invention does not require that all network documents be filtered for a particular user, as does the prior art.

Client **20** browses the web aided by Personal Web **12**. In browsing mode, the contents of a web site are customized according to the User Model **13**. Personal Web interacts with a cooperating web site by supplying User Model information, and a web page authored in a dynamic language (e.g., DHTML) is personalized to the user's profile. In navigation mode, a personal navigation aid suggests to the user relevant

links within the visited site or outside it given the context, for example, the current web page and previously visited pages, and knowledge of the user profile.

Client **22** illustrates the find-an-expert mode of Personal Web **12**. The user supplies an expert information or product need in the form of a sample web page or text string, and Personal Web **12** locates an expert in the user's company, circle of friends, or outside groups that has the relevant information and expertise, based on the expert's User Model **13**. The located expert not only has the correct information, but presents it in a manner of most interest to the user, for example, focussing on technical rather than business details of a product.

Client **24** uses the personal pushed information mode of Personal Web **12**. Personal Web **12** collects and presents personal information to a user based on the User Model **13**. The pushed information is not limited to a fixed or category or topic, but includes any information of interest to the user. In communities, organizations, or group of users, the pushed information can include automatic routing and delivery of newly created documents that are relevant to the users.

Finally, client **26** illustrates the product recommendation mode of Personal Web **12**. The user submits a query for information about a product type, and Personal Web **12** locates the products and related information that are most relevant to the user, based on the User Model **13**. As described further below, product information is gathered from all available knowledge sources, such as product reviews and press releases, and Personal Web **12** can recommend a product that has never been purchased or rated by any users.

All of the above features of Personal Web **12** are based on a User Model **13** that represents user interests in a document or product independently of any specific user information need, i.e., not related to a specific query. The User Model **13** is a function that is developed and updated using a variety of knowledge sources and that is independent of a specific representation or data structure. The underlying mathematical framework of the modeling and training algorithms discussed below is based on Bayesian statistics, and in particular on the optimization criterion of maximizing posterior probabilities. In this approach, the User Model is updated based on both positive and negative training examples. For example, a search result at the top of the list that is not visited by the user is a negative training example.

The User Model **13**, with its associated representations, is an implementation of a learning machine. As defined in the art, a learning machine contains tunable parameters that are altered based on past experience. Personal Web **12** stores parameters that define a User Model **13** for each user, and the parameters are continually updated based on monitored user interactions while the user is engaged in normal use of a computer. While a specific embodiment of the learning machine is discussed below, it is to be understood that any model that is a learning machine is within the scope of the present invention.

The present invention can be considered to operate in three different modes: initialization, updating or dynamic learning, and application. In the initialization mode, a User Model **13** is developed or trained based in part on a set of user-specific documents. The remaining two modes are illustrated in the block diagram of FIG. **2**. While the user is engaged in normal use of a computer, Personal Web **12** operates in the dynamic learning mode to transparently monitor user interactions with data (step **30**) and update the User Model **13** to reflect the user's current interests and needs. This updating is performed by updating a set of user-specific data files in step **32**, and then using the data files to update the parameters of the User Model

13 in step 34. The user-specific data files include a set of documents and products associated with the user, and monitored user interactions with data. Finally, Personal Web 12 applies the User Model 13 to unseen documents, which are first analyzed in step 36, to determine the user's interest in the document (step 38), and performs a variety of services based on the predicted user interest (step 40). In response to the services provided, the user performs a series of actions, and these actions are in turn monitored to further update the User Model 13.

The following notation is used in describing the present invention. The user and his or her associated representation are denoted with u , a user query with q , a document with d , a product or service with p , a web site with s , topic with t , and a term, meaning a word or phrase, with w . The term "document" includes not just text, but any type of media, including, but not limited to, hypertext, database, spreadsheet, image, sound, and video. A single document may have one or multiple distinct media types. Accordingly, the set of all possible documents is D , the set of all users and groups is U , the set of all products and services is P , etc. The user information or product need is a subset of D or P . Probability is denoted with P , and a cluster of users or of clusters with c , with which function semantics are used. For example, $c(c(u))$ is the cluster of clusters in which the user u is a member ("the grandfather cluster"). Note that an explicit notation of world knowledge, such as dictionaries, atlases, and other general knowledge sources, which can be used to estimate the various posterior probabilities, is omitted.

A document classifier is a function whose domain is any document, as defined above, and whose range is the continuous interval $[0,1]$. For example, a document classifier may be a probability that a document d is of interest to a particular user or a group of users. Specific document classifiers of the present invention are obtained using the User Model 13 and Group Model. The User Model 13 represents the user interest in a document independent of any specific user information need. This estimation is unique to each user. In strict mathematical terms, given a user u and a document d , the User Model 13 estimates the probability $P(uld)$. $P(uld)$ is the probability of the event that the user u is interested in the document d , given everything that is known about the document d . This classifier is extended to include $P(uld,con)$, the probability that a user is interested in a given document based on a user's current context, for example, the web pages visited during a current interaction session.

The Group or Cluster Model is a function that represents the interest level of a group of users in a document independently of any specific information need. For example, for the group of users $c(u)$, the mathematical notation of this probability, which is determined by applying the Group Model to a document d , is $P(c(u)|d)$.

A schematic diagram of the User Model is shown in FIG. 3, which illustrates the various knowledge sources (in circles) used as input to the User Model. The knowledge sources are used to initialize and update the User Model, so that it can accurately take documents and generate values of user interest in the documents, given the context of the user interaction. Note that some of the knowledge sources are at the individual user level, while others refer to aggregated data from a group of users, while still others are independent of all users. Also illustrated in FIG. 3 is the ability of the User Model to estimate a user interest in a given product, represented mathematically as the interest of a user in a particular document, given that the document describes the product: $P(\text{user|document, product described}=p)$. As explained further below, the long-term user interest in a product is one of many

probabilities incorporated into the computation of user interest in all documents, but it can also be incorporated into estimation of a current user interest in a product.

Beginning at the bottom left of FIG. 3, User Data and Actions include all user-dependent inputs to the User Model, including user browser documents, user-supplied documents, other user-supplied data, and user actions, such as browsing, searching, shopping, finding experts, and reading news. Data and actions of similar users are also incorporated into the User Model by clustering all users into a tree of clusters. Clustering users allows estimation of user interests based on the interests of users similar to the user. For example, if the user suddenly searches for information in an area that is new to him or her, the User Model borrows characteristics of User Models of users with similar interests. Topic classifiers are used to classify documents automatically into topics according to a predefined topic tree. Similarly, product models determine the product or product categories, if any, referred to by a document. Product models also extract relevant feature of products from product-related documents. The topic experts input provides input of users with a high interest in a particular topic, as measured by their individual User Models. Finally, the User Model incorporates world knowledge sources that are independent of all users, such as databases of company names, yellow pages, thesauri, dictionaries, and atlases.

User Model Representations

Given the inputs shown in FIG. 3, the User Model is a function that may be implemented with any desired data structure and that is not tied to any specific data structure or representation. The following currently preferred embodiment of abstract data structures that represent the User Model 13 is intended to illustrate, but not limit, the User Model of the present invention. Some of the structures hold data and knowledge at the level of individual users, while others store aggregated data for a group or cluster of users. Initialization of the various data structures of the User Model is described in the following section; the description below is of the structures themselves.

User-dependent inputs are represented by components of the User Model shown in FIGS. 4A-4E. These inputs are shown as tables for illustration purposes, but may be any suitable data structure. The user-dependent components include an informative word or phrase list, a web site distribution, a user topic distribution, a user product distribution, and a user product feature distribution. Each of these user-dependent data structures can be thought of as a vector of most informative or most frequent instances, along with a measure representing its importance to the user.

The informative word and phrase list of FIG. 4A contains the most informative words and phrases found in user documents, along with a measure of each informative phrase or word's importance to the user. As used herein, an "informative phrase" includes groups of words that are not contiguous, but that appear together within a window of a predefined number of words. For example, if a user is interested in the 1999 Melissa computer virus, then the informative phrase might include the words "virus," "Melissa," "security," and "IT," all appearing within a window of 50 words. The sentence "The computer virus Melissa changed the security policy of many IT departments" corresponds to this phrase.

In addition to the words and phrases, the list contains the last access time of a document containing each word or phrase and the total number of accessed documents containing the words. One embodiment of the informative measure is a word probability distribution $P(w|u)$ representing the interest of a user u in a word or phrase w , as measured by the word's

frequency in user documents. Preferably, however, the informative measure is not simply a measure of the word frequency in user documents; common words found in many documents, such as "Internet," provide little information about the particular user's interest. Rather, the informative measure should be high for words that do not appear frequently across the entire set of documents, but whose appearance indicates a strong likelihood of the user's interest in a document. A preferred embodiment uses the TFIDF measure, described in Ricardo Baeza-Yates and Berthier Ribeiro-Neto, *Modern Information Retrieval*, Addison Wesley, 1999, in which TF stands for term frequency, and IDF stands for inverse document frequency. Mathematically, if $f_{u,w}$ denotes the frequency of the word w in user u documents, and D_w denotes the number of documents containing the word w , then the importance of a word w to a user u is proportional to the product $f_{u,w} \cdot D/D_w$.

A more preferred embodiment of the measure of each word's importance uses a mathematically sound and novel implementation based on information theory principles. In particular, the measure used is the mutual information between two random variables representing the user and the word or phrase. Mutual information is a measure of the amount of information one random variable contains about another; a high degree of mutual information between two random variables implies that knowledge of one random variable reduces the uncertainty in the other random variable.

For the present invention, the concept of mutual information is adapted to apply to probability distributions on words and documents. Assume that there is a document in which the user's interest must be ascertained. The following two questions can be asked: Does the phrase p appear in the document?; and Is the document of interest to the user u ? Intuitively, knowing the answer to one of the questions reduces the uncertainty in answering the other question. That is, if the word w appears in a different frequency in the documents associated with the user u from its frequency in other documents, it helps reduce the uncertainty in determining the interest of user u in the document.

Through the concept of mutual information, information theory provides the mathematical tools to quantify this intuition in a sound way. For a detailed explanation, see T. Cover and J. Thomas, *Elements of Information Theory*, Wiley, 1991. In this embodiment of the informative measure, two indicator variables are defined. I_w has a value of 1 when the word w appears in a web document and 0 when it does not, and I_u has a value of 1 when a web document is of interest to the user u and 0 when it does not. The mutual information between the two random variables I_w and I_u is defined as:

$$I(I_w; I_u) = \sum_{i_w \in \mathcal{I}_w} \sum_{i_u \in \mathcal{I}_u} P(i_w, i_u) \log_2 \frac{P(i_w, i_u)}{P(i_w)P(i_u)}$$

The probabilities in this formula are computed over a set of documents of interest to the user and a set of documents not of interest to the user. For example, consider a set of 100 documents of interest to the user, and a set of 900 documents not of interest to the user. Then $P(i_u=1)=0.1$, and $P(i_u=0)=0.9$. Assume that in the combined set of 1000 documents, 150 contain the word "Bob." Then $P(i_w=1)=0.15$, and $P(i_w=0)=0.85$. In addition, assume that "Bob" appears in all 100 of the documents of interest to the user. $P(i_w, i_u)$ has the following four values:

i_u	i_w	$P(i_w, i_u)$
0	0	850/1000
0	1	50/1000
1	0	0/1000
1	1	100/1000

Using the above formula, the mutual information between the user and word Bob is:

$$\begin{aligned} I(I_{Bob}; I_{user}) &= 850/1000 \log [850/1000 / (0.85 * 0.9)] + \\ &50/1000 \log [50/1000 / (0.15 * 0.9)] + \\ &0/1000 \log [0/1000 / (0.1 * 0.85)] + \\ &100/1000 \log [100/1000 / (0.15 * 0.1)] \\ &= 0.16. \end{aligned}$$

Mutual information is a preferred measure for selecting the word and phrase list for each user. The chosen words and phrases have the highest mutual information.

The remaining User Model representations are analogously defined using probability distributions or mutual information. The web site distribution of FIG. 4B contains a list of web sites favored by the user along with a measure of the importance of each site. Given the dynamic nature of the Internet, in which individual documents are constantly being added and deleted, a site is defined through the first backslash (after the www). For example, the uniform resource locator (URL) <http://www.herring.com/companies/2000> . . . is considered as www.herring.com. Sites are truncated unless a specific area within a site is considered a separate site; for example, www.cnn.com/health is considered to be a different site than www.cnn.com/us. Such special cases are decided experimentally based on the amount of data available on each site and the principles of data-driven approaches, described in Vladimir S. Cherkassky and Filip M. Mulier, *Learning from Data: Concepts, Theory, and Methods*, in *Adaptive and Learning Systems for Signal Processing, Communications and Control*, Simon Haykin, series editor, Wiley & Sons, March, 1998. Each site has an importance measure, either a discrete probability distribution, $P(s|u)$, representing the interest of user u in a web site s , or the mutual information metric defined above, $I(I_s; I_u)$, representing the mutual information between the user u and a site s . The web site distribution also contains the last access time and number of accesses for each site.

FIG. 4C illustrates the user topic distribution, which represents the interests of the user in various topics. The user topic distribution is determined from a hierarchical, user-independent topic model, for example a topic tree such as the Yahoo directory or the Open Directory Project, available at <http://dmoz.org/>. Each entry in the tree has the following form:

Computers\Internet\WWW\Searching the Web\Directories\Open Directory Project\

where the topic following a backslash is a child node of the topic preceding the backslash. The topic model is discussed in more detail below.

For each node of the topic tree, a probability is defined that specifies the user interest in the topic. Each level of the topic

13

model is treated distinctly. For example, for the top level of the topic model, there is a distribution in which

$$P(t_u|u)+P(t_1|u)=1,$$

where t_1 represents the top level of topics and is the same set of topics for each user, e.g., technology, business, health, etc. $P(t_1|u)$ is the sum of the user probabilities on all top level topics. For each topic level, t_u represents specific interests of each user that are not part of any common interest topics, for instance family and friends' home pages. For lower topic levels, every node in the tree is represented in the user topic distribution by a conditional probability distribution. For example, if the Technology node splits into Internet, Communication, and Semiconductors, then the probability distribution is of the form:

$$P(\text{Internet}|u, \text{Technology})+P(\text{Communication}|u, \text{Technology})+P(\text{Semiconductors}|u, \text{Technology})+P(t_u|u, \text{Technology})=1$$

Rather than probabilities, the mutual information metric defined above may be used; $I(I_t; I_u)$ represents the mutual information between the user u and the topic t . An exemplary data structure shown in FIG. 4C for storing the user topic distribution contains, for each topic, the topic parent node, informative measure, last access time of documents classified into the topic, and number of accesses of documents classified into the topic. Note that the User Model contains an entry for every topic in the tree, some of which have a user probability or mutual information of zero.

The user product distribution of FIG. 4D represents the interests of the user in various products, organized in a hierarchical, user-independent structure such as a tree, in which individual products are located at the leaf nodes of the tree. The product taxonomy is described in further detail below. The product taxonomy is similar to the topic tree. Each entry in the tree has the following form:

Consumer Electronics\Cameras\Webcams\3Com Home-Connect\

where a product or product category following a backslash is a child node of a product category preceding the backslash.

For each node of the product model, a probability is defined that specifies the user interest in that particular product or product category. Each level of the product model is treated distinctly. For example, for the top level of the product hierarchy, there is a distribution in which

$$P(p_1|u)=1,$$

where p_1 represents the top level of product categories and is the same for each user, e.g., consumer electronics, computers, software, etc. For lower product category levels, every node in the tree is represented in the user product distribution by a conditional probability distribution. For example, if the Cameras node splits into Webcams and Digital Cameras, then the probability distribution is of the form:

$$P(\text{Webcams}|u, \text{Cameras})+P(\text{Digital Cameras}|u, \text{Cameras})=1$$

Rather than probabilities, the mutual information metric defined above may be used. Then $I(I_p; I_u)$ represents the mutual information between the user u and the product or product category p . An exemplary data structure for storing the user product distribution contains, for each product, the product ID, product parent node, user probability, last purchase time of the product, number of product purchases, last access time of documents related to the product, and number of related documents accessed.

14

For each product or category on which the user has a nonzero probability, the User Model contains a user product feature distribution on the relevant features, as shown in FIG. 4E. Each product category has associated with it a list of features, and the particular values relevant to the user are stored along with a measure of the value's importance, such as a probability $P(f|u, p)$ or mutual information measure $I(I_p; I_u)$. For example, Webcams have a feature Interface with possible values Ethernet (10BaseT), Parallel, PC Card, serial, USB, and TV. Probability values of each feature sum to one; that is,

$$P(\text{Ethernet}|u, \text{Interface, Webcam})+P(\text{Parallel}|u, \text{Interface, Webcam})+P(\text{PC Card}|u, \text{Interface, Webcam})+P(\text{serial}|u, \text{Interface, Webcam})+P(\text{USB}|u, \text{Interface, Webcam})+P(\text{TV}|u, \text{Interface, Webcam})=1.$$

User probability distributions or mutual information measures are stored for each feature value of each node. Note that there is no user feature value distribution at the leaf nodes, since specific products have particular values of each feature.

Finally, user-dependent components of the User Model include clusters of users similar to the user. Users are clustered into groups, forming a cluster tree. One embodiment of a user cluster tree, shown in FIG. 5A, hard classifies users into clusters that are further clustered. Each user is a member of one and only one cluster. For example, Bob is clustered into a cluster $c(u)$, which is further clustered into clusters of clusters, until the top level cluster is reached $c(U)$. The identity of the user's parent cluster and grandfather cluster is stored as shown in FIG. 5B, and information about the parent cluster is used as input into the User Model. As described below, clusters are computed directly from User Models, and thus need not have a predefined semantic underpinning.

Preferably, the User Model does not use hard clustering, but rather uses soft or fuzzy clustering, also known as probabilistic clustering, in which the user belongs to more than one cluster according to a user cluster distribution $P(c(u))$. FIG. 6A illustrates fuzzy clusters in a cluster hierarchy. In this case, Bob belongs to four different clusters according to the probability distribution shown. Thus Bob is most like the members of cluster C4, but still quite similar to members of clusters C1, C2, C3, and C4. Fuzzy clustering is useful for capturing different interests of a user. For example, a user may be a small business owner, a parent of a small child, and also an avid mountain biker, and therefore need information for all three roles. Probabilistic clustering is described in detail in the Ph.D. thesis of Steven J. Nowlan, "Soft Competitive Adaptation: Neural Network Learning Algorithms Based on Fitting Statistical Mixtures," School of Computer Science, Carnegie Mellon University, Pittsburgh, Pa., 1991. A suitable data structure for representing fuzzy clusters is shown in FIG. 6B. Each row stores the cluster or user ID, one parent ID, and the cluster probability, a measure of similarity between the cluster or user and the parent cluster.

Note that all elements of an individual User Model for a user u also apply to a cluster of users $c(u)$. Thus for each cluster, a Group Model is stored containing an informative word list, a site distribution, a topic distribution, a group product distribution, and a group product feature distribution, each with appropriate measures. For example, $P(p|c(u))$ represents the interest of a cluster $c(u)$ in various products p .

The user-dependent User Model representations also include a user general information table, which records global information describing the user, such as the User ID, the number of global accesses, the number of accesses within a recent time period, and pointers to all user data structures.

Other knowledge sources of the User Model are independent of the user and all other users. Topic classifiers are used

to classify documents into topics according to a predefined topic tree, an example of which is illustrated in FIG. 7. A variety of topic trees are available on the web, such as the Yahoo directory or Open Directory Project (www.dmoz.org). A topic classifier is a model similar to the user model that estimates the probability that a document belongs to a topic. Every node on the topic tree has a stored topic classifier. Thus the set of all topic classifiers computes a probability distribution of all of the documents in the set of documents D among the topic nodes. For example, the topic classifier in the root node in FIG. 7 estimates the posterior probabilities $P(t|d)$, where t represents the topic of document d and is assigned values from the set {Arts, Business, Health, News, Science, Society}. Similarly, the topic classifier for the Business node estimates the posterior probability $P(t|d, \text{Business})$, where t represents the specific topic of the document d within the Business category. Mathematically, this posterior probability is denoted $P(t|d)=\text{Business}\backslash\text{Investing}\backslash t(d)=\text{Business}, d)$, which represents the probability that the subtopic of the document d within Business is Investing, given that the topic is Business. The topic tree is stored as shown in FIG. 8, a table containing, for each node, the topic ID, depth level, topic parent ID, number of child nodes, and topic ID of the child nodes.

The topic experts model estimates the probability that a document is of interest to users who are interested in a particular topic, independent of any specific user information need. Each node of the topic tree has, in addition to a topic classifier, a corresponding topic expert function. Note that the topic classifier and topic expert function are independent; two documents can be about investing, but one of high interest to expert users and the other of no interest to expert users. The topic expert model can be considered an evaluation of the quality of information in a given document. The assumption behind the topic experts model is that the degree of interest of a user in a given topic is his or her weight for predicting the quality or general interest level in a document classified within the particular topic. Obviously there are outliers to this assumption, for example, novice users. However, in general and averaged across many users, this measure is a good indicator of a general interest level in a document. For every topic in the tree, a list of the N clusters with the most interest in the topic based on the cluster topic distribution is stored. The cluster topic distribution is similar to the user topic distribution described above, but is averaged over all users in the cluster. An exemplary data structure for storing the topic experts model is shown in FIG. 9.

Finally, a product model is stored for every node of a product taxonomy tree, illustrated in FIG. 10. Examples of product taxonomy trees can be found at www.cnet.com and www.productopia.com, among other locations. In any product taxonomy tree, the leaf nodes, i.e., the bottom nodes of the tree, correspond to particular products, while higher nodes represent product categories. Product models are similar to topic classifiers and User Models, and are used to determine whether a document is relevant to a particular product or product category. Thus a product model contains a list of informative words, topics, and sites. The set of all product models computes a probability distribution of all of the documents in the set of documents D among the product nodes. For example, the product model in the root node in FIG. 10 estimates the posterior probabilities $P(p|d)$, where p represents the product referred to in document d and is assigned values from the set {Consumer Electronics, Computers, Software}. Similarly, the product model for the Consumer Electronics node estimates the posterior probability $P(p|d, \text{Consumer Electronics})$, where p represents the product category

of the document d within the Consumer Electronics category. Mathematically, this posterior probability is denoted $P(p(d)=\text{Consumer Electronics}\backslash\text{CD Players}\backslash p(d)=\text{Consumer Electronics}, d)$, which represents the probability that the subproduct category of the document d within Consumer Electronics is CD Players, given that the product category is Consumer Electronics. The product tree is stored as shown in FIG. 11, a table containing, for each node, the topic ID, depth level, topic parent ID, number of child nodes, and topic ID of the child nodes.

Each node of the product tree has an associated product feature list, which contains particular descriptive features relevant to the product or category. Nodes may have associated feature values; leaf nodes, which represent specific products, have values of all relevant product features. Product feature lists are determined by a human with knowledge of the domain. However, feature values may be determined automatically from relevant knowledge sources as explained below.

For example, in the product tree of FIG. 10, CD Players is the parent node of the particular CD player Sony CDP-CX350 and Harman Kardon CDR2. The product category CD Players has the following features: Brand, CD Capacity, Digital Output, Plays Minidisc, and Price Range. Each feature has a finite number of potential feature values; for example, CD Capacity has potential feature values 1 Disc, 1-10 Discs, 10-50 Discs, or 50 Discs or Greater. Individual products, the child nodes of CD Players, have one value of each feature. For example, the Sony CDP-CX350 has a 300 disc capacity, and thus a feature value of 50 Discs or Greater.

Some product features are relevant to multiple product categories. In this case, product features propagate as high up the product tree as possible. For example, digital cameras have the following product features: PC Compatibility, Macintosh Compatibility, Interfaces, Viewfinder Type, and Price Range. Webcams have the following product features: PC Compatibility, Macintosh Compatibility, Interfaces, Maximum Frames per Second, and Price Range. Common features are stored at the highest possible node of the tree; thus features PC Compatibility, Macintosh Compatibility, and Interfaces are stored at the Cameras node. The Digital Cameras node stores only product feature Viewfinder Type, and the Webcams node stores only product feature Maximum Frames per Second. Note that product feature Price Range is common to CD Players and Cameras, and also Personal Minidiscs, and thus is propagated up the tree and stored at node Consumer Electronics.

Individual products at leaf nodes inherit relevant features from all of their ancestor nodes. For example, Kodak CD280 inherits the feature Viewfinder Type from its parent; PC Compatibility, Macintosh Compatibility, and Interfaces from its grandparent; and Price Range from its great-grandparent. A product feature list is stored as shown in FIG. 12A, and contains, for each product ID, the associated feature and its value. All potential feature values are stored in a product feature value list, as shown in FIG. 12B.

The system also includes a document database that indexes all documents D . The document database records, for each document, a document ID, the full location (the URL of the document), a pointer to data extracted from the document, and the last access time of the document by any user. A word database contains statistics of each word or phrase from all user documents. The word database contains the word ID, full word, and word frequency in all documents D , used in calculating informative measures for individual users and clusters.

Initialization of User Model

The User Model is initialized offline using characterizations of user behavior and/or a set of documents associated with the user. Each data structure described above is created during initialization. In other words, the relevant parameters of the learning machine are determined during initialization, and then continually updated online during the update mode.

In one embodiment, the user documents for initializing the User Model are identified by the user's web browser. Most browsers contain files that store user information and are used to minimize network access. In Internet Explorer, these files are known as favorites, cache, and history files. Most commercial browsers, such as Netscape Navigator, have equivalent functionality; for example, bookmarks are equivalent to favorites. Users denote frequently-accessed documents as bookmarks, allowing them to be retrieved simply by selection from the list of bookmarks. The bookmarks file includes for each listing its creation time, last modification time, last visit time, and other information. Bookmarks of documents that have changed since the last user access are preferably deleted from the set of user documents. The Internet Temporary folder contains all of the web pages that the user has opened recently (e.g., within the last 30 days). When a user views a web page, it is copied to this folder and recorded in the cache file, which contains the following fields: location (URL), first access time, and last access time (most recent retrieval from cache). Finally, the history file contains links to all pages that the user has opened within a set time period.

Alternatively, the user supplies a set of documents, not included in any browser files, that represent his or her interests. The User Model can also be initialized from information provided directly by the user. Users may fill out forms, answer questions, or play games that ascertain user interests and preferences. The user may also rate his or her interest in a set of documents provided.

User documents are analyzed as shown in FIG. 13 to determine initial parameters for the various functions of the User Model. A similar analysis is used during updating of the User Model. Note that during updating, both documents that are of interest to the user and documents that are not of interest to the user are analyzed and incorporated into the User Model. The process is as follows. In a first step 82, the format of documents 80 is identified. In step 84, documents 80 are parsed and separated into text, images and other non-text media 88, and formatting. Further processing is applied to the text, such as stemming and tokenization to obtain a set of words and phrases 86, and information extraction. Through information extraction, links 90 to other documents, email addresses, monetary sums, people's names, and company names are obtained. Processing is performed using natural language processing tools such as LinguistX® and keyword extraction tools such as Thing Finder™, both produced by Inxight (www.inxight.com). Further information on processing techniques can be found in Christopher D. Manning and Hinrich Schutze, *Foundations of Statistical Natural Language Processing*, MIT Press, 1999. Additional processing is applied to images and other non-text media 88. For example, pattern recognition software determines the content of images, and audio or speech recognition software determines the content of audio. Finally, document locations 94 are obtained.

Parsed portions of the documents and extracted information are processed to initialize or update the user representations in the User Model. In step 96, user informative words or phrases 98 are obtained from document words and phrases 86. In one embodiment, a frequency distribution is obtained to calculate a TFIDF measure quantifying user interest in words 98. Alternatively, mutual information is calculated between the two indicator variables I_w and I_u as explained

above. The set of informative words 98 contains words with the highest probabilities or mutual information.

In step 100, the topic classifiers are applied to all extracted information and portions of documents 80 to obtain a probability distribution $P(t|d)$ for each document on each node of the topic tree. As a result, each node has a set of probabilities, one for each document, which is averaged to obtain an overall topic node probability. The average probabilities become the initial user topic distribution 102. If desired, mutual information between the two indicator variables I_t and I_u can be determined as explained above.

Similarly, in step 104, product models are applied to all extracted information from documents 80 to classify documents according to the product taxonomy tree. From user purchase history 105, additional product probabilities are obtained. Probabilities for each node are combined, weighting purchases and product-related documents appropriately, to obtain a user product distribution 106. Note that only some of documents 80 contain product-relevant information and are used to determine the user product distribution 106. Product models return probabilities of zero for documents that are not product related.

The user product feature distribution 108 can be obtained from different sources. If a user has a nonzero probability for a particular product node, then the feature distribution on that node is obtained from its leaf nodes. For example, if one of the user documents was classified into Kodak DC280 and another into Nikon Coolpix 950, then the user product feature distribution for the Digital Cameras node has a probability of 0.5 for the feature values corresponding to each camera. Feature value distributions propagate throughout the user product feature distributions. For example, if the two cameras are in the same price range, \$300-\$400, then the probability of the value \$300-\$400 of the feature Price Range is 1.0, which propagates up to the Consumer Electronics node (assuming that the user has no other product-related documents falling within Consumer Electronics).

Alternatively, product feature value distributions are obtained only from products that the user has purchased, and not from product-related documents in the set of user documents. Relevant feature values are distributed as high up the tree as appropriate. If the user has not purchased a product characterized by a particular feature, then that feature has a zero probability. Alternatively, the user may explicitly specify his or her preferred feature values for each product category in the user product distribution. User-supplied information may also be combined with feature value distributions obtained from documents or purchases.

Document locations 94 are analyzed (step 110) to obtain the user site distribution 112. Analysis takes into account the relative frequency of access of the sites within a recent time period, weighted by factors including how recently a site was accessed, whether it was kept in the favorites or bookmarks file, and the number of different pages from a single site that were accessed. Values of weighting factors are optimized experimentally using jackknifing and cross-validation techniques described in H. Bourlard and N. Morgan, *Connectionist Speech Recognition: A Hybrid Approach*, Kluwer Academic Publishers, 1994.

Note that there is typically overlap among the different representations of the User Model. For example, a news document announcing the release of a new generation of Microsoft servers has relevant words Microsoft and server. In addition, it is categorized within the product taxonomy under Microsoft servers and the topic taxonomy under computer hardware. This document may affect the user's word list, product distribution, and topic distribution. After the User

Models are initialized for all users, cluster membership can be obtained. Clusters contain users with a high degree of similarity of interests and information needs. A large number of clustering algorithms are available; for examples, see K. Fukunaga, *Statistical Pattern Recognition*, Academic Press, 1990. As discussed above, users are preferably soft clustered into more than one cluster. Preferably, the present invention uses an algorithm based on the relative entropy measure from information theory, a measure of the distance between two probability distributions on the same event space, described in T. Cover and J. Thomas, *Elements of Information Theory*, Chapter 2, Wiley, 1991. Clustering is unsupervised. That is, clusters have no inherent semantic significance; while a cluster might contain users with a high interest in mountain biking, the cluster tree has no knowledge of this fact.

In a preferred embodiment, the relative entropy between two User Model distributions on a fixed set of documents D_{sample} is calculated. D_{sample} is chosen as a good representation of the set of all documents D . Distributions of similar users have low relative entropy, and all pairs of users within a cluster have relative entropy below a threshold value. The User Model of each user is applied to the documents to obtain a probability of interest of each user in each document in the set. The relative entropy between two user distributions for a single document is calculated for each document in the set, and then summed across all documents.

The exact mathematical computation of the relative entropy between two users is as follows. An indicator variable $I_{u,d}$ is assigned to 1 when a document d is of interest to a user u and 0 when it is not. For two users u_1 and u_2 and for any document d , the relative entropy between the corresponding distributions is:

$$D(I_{u_1,d} \parallel I_{u_2,d}) = \sum_{i \in I} P(i_{u_1,d}) \log_2 \frac{P(i_{u_1,d})}{P(i_{u_2,d})}$$

For example, if $P(u_1|d)=0.6$ and $P(u_2|d)=0.9$, then

$$D(I_{u_1,d} \parallel I_{u_2,d}) = 0.4 \log(0.4/0.1) + 0.6 \log(0.6/0.9).$$

The relative entropy can be converted to a metric D' that obeys the triangle inequality:

$$D'(I_1 \parallel I_2) = 0.5 * (D(I_1 \parallel I_2) + D(I_2 \parallel I_1)).$$

For any two users u_1 and u_2 , and for each document in D_{sample} , the metric D' is computed between the corresponding indicator variable distributions on the document. The values for all document are summed, and this sum is the distance metric for clustering users. This distance is defined as:

$$\text{Distance}(u_1, u_2) = \sum_{d_j \in D_{sample}} D'(I_{u_1,d_j} \parallel I_{u_2,d_j}),$$

An alternative clustering algorithm computes the relative entropy between individual user distributions in the User Model, for example, between all informative word lists, site distributions, etc., of each user. The equations are similar to those above, but compute relative entropy based on indicator variables such as $I_{u,w}$, which is assigned a value of 1 when a word w is of interest to a user u . The calculated distances between individual user distributions on words, sites, topics, and products are summed to get an overall user distance. This second algorithm is significantly less computationally costly than the preferred algorithm above; selection of an algorithm

depends on available computing resources. In either case, relative entropy can also be computed between a user and cluster of users.

Each cluster has a Group or Cluster Model that is analogous to a User Model. Cluster Models are generated by averaging each component of its members' User Models. When fuzzy clusters are used, components are weighted by a user's probability of membership in the cluster.

In some cases, initialization is performed without any user-specific information. A user may not have a large bookmarks file or cache, or may not want to disclose any personal information. For such users, prototype users are supplied. A user can choose one or a combination of several prototype User Models, such as the technologist, the art lover, and the sports fan. Predetermined parameters of the selected prototype user are used to initialize the User Model. Users can also opt to add only some parameters of a prototype user to his or her existing User Model by choosing the prototype user's distribution of topics, words, sites, etc. Note that prototype users, unlike clusters, are semantically meaningful. That is, prototype users are trained on a set of documents selected to represent a particular interest. For this reason, prototype users are known as "hats," as the user is trying on the hat of a prototype user.

Users can also choose profiles on a temporary basis, for a particular session only. For example, in a search for a birthday present for his or her teenage daughter, a venture capitalist from Menlo Park may be interested in information most probably offered to teenagers, and hence may choose a teenage girl profile for the search session.

User-independent components are also initialized. The topic classifiers are trained using the set of all possible documents D . For example, D may be the documents classified by the Open Directory Project into its topic tree. Topic classifiers are similar to a User Model, but with a unimodal topic distribution function (i.e., a topic model has a topic distribution value of 1 for itself and 0 for all other topic nodes). The set of documents associated with each leaf node of the topic tree is parsed and analyzed as with the user model to obtain an informative word list and site distribution. When a topic classifier is applied to a new document, the document's words and location are compared with the informative components of the topic classifier to obtain $P(t|d)$. This process is further explained below with reference to computation of $P(u|d)$. Preferably, intermediate nodes of the tree do not have associated word list and site distributions. Rather, the measures for the word list and site distribution of child nodes are used as input to the topic classifier of their parent nodes. For example, the topic classifier for the Business node of the topic tree of FIG. 7 has as its input the score of the site of the document to be classified according to the site distributions of the topic models of its child nodes, Employment, Industries, and Investing. The classifier can be any non-linear classifier such as one obtained by training a Multilayer Perceptron (MLP) using jackknifing and cross-validation techniques, as described in H. Bourlard and N. Morgan, *Connectionist Speech Recognition: A Hybrid Approach*, Kluwer Academic Publishers, 1994. It can be shown that a MLP can be trained to estimate posterior probabilities; for details, see J. Hertz, A. Krogh, R. Palmer, *Introduction to The Theory of Neural Computation*, Addison-Wesley, 1991.

The topic experts model is initialized by locating for every node in the topic tree the N clusters that are of the same depth in the user cluster tree as the user, and that have the highest interest in the topic, based on their cluster topic distribution. The cluster topic distribution $P(t|c(u))$ is simply an average of the user topic distribution $P(t|u)$ for each user in the cluster. The topic experts model is used to determine the joint prob-

ability that a document and the topic under consideration are of interest to any user, $P(t,d)$. Using Bayes' rule, this term can be approximated by considering the users of the N most relevant clusters.

$$P(t, d) = \sum_{i \in N} P(c_i | t, d) P(t | d) P(d)$$

The topic experts model is, therefore, not a distinct model, but rather an ad hoc combination of user and cluster topic distributions and topic models.

Product models are initialized similarly to User Models and topic classifiers. Each leaf node in the product tree of FIG. 10 has an associated set of documents that have been manually classified according to the product taxonomy. These documents are used to train the product model as shown for the User Model in FIG. 13. As a result, each leaf node of the product tree contains a set of informative words, a topic distribution, and a site distribution. Each node also contains a list of features relevant to that product, which is determined manually. From the documents, values of the relevant features are extracted automatically using information extraction techniques to initialize the feature value list for the product. For example, the value of the CD Capacity is extracted from the document. Information extraction is performed on unstructured text, such as HTML documents, semi-structured text, such as XML documents, and structured text, such as database tables. As with the topic model, a nonlinear function such as a Multilayer Perceptron is used to train the product model.

Preferably, as for topic classifiers, intermediate nodes of the product tree do not have associated word lists, site distributions, and topic distributions. Rather, the measures for the word list, site distribution, and topic distribution of child nodes are used as input to the product models of their parent nodes. Alternatively, each parent node may be trained using the union of all documents of its child nodes.

Updating the User Model

The User Model is a dynamic entity that is refined and updated based on all user actions. User interactions with network data are transparently monitored while the user is engaged in normal use of his or her computer. Multiple distinct modes of interaction of the user are monitored, including network searching, network navigation, network browsing, email reading, email writing, document writing, viewing pushed information, finding expert advice, product information searching, and product purchasing. As a result of the interactions, the set of user documents and the parameters of each user representation in the User Model are modified.

While any nonlinear function may be used in the User Model (e.g., a Multilayer Perceptron), a key feature of the model is that the parameters are updated based on actual user reactions to documents. The difference between the predicted user interest in a document or product and the actual user interest becomes the optimization criterion for training the model.

Through his or her actions, the user creates positive and negative patterns. Positive examples are documents of interest to a user: search results that are visited following a search query, documents saved in the user favorites or bookmarks file, web sites that the user visits independently of search queries, etc. Negative examples are documents that are not of interest to the user, and include search results that are ignored although appear at the top of the search result, deleted book-

marks, and ignored pushed news or email. Conceptually, positive and negative examples can be viewed as additions to and subtractions from the user data and resources.

Information about each document that the user views is stored in a recently accessed buffer for subsequent analysis. The recently accessed buffer includes information about the document itself and information about the user's interaction with the document. One possible implementation of a buffer is illustrated in FIG. 14; however, any suitable data structure may be used. The recently-accessed buffer contains, for each viewed document, a document identifier (e.g., its URL); the access time of the user interaction with the document; the interaction type, such as search or navigation; the context, such as the search query; and the degree of interest, for example, whether it was positive or negative, saved in the bookmarks file, how long the user spent viewing the document, or whether the user followed any links in the document. Additional information is recorded for different modes of interaction with a document as discussed below.

A metric is determined for each document to indicate whether it is a positive, negative or neutral event; this metric can potentially be any grade between 0 and 1, where 0 is a completely negative event, 1 is a completely positive event, and 0.5 is a neutral event. Previous user interactions may be considered in computing the metric; for example, a web site that the user accesses at a frequency greater than a predetermined threshold frequency is a positive example. After each addition to or subtraction from the set of user documents, the document is parsed and analyzed as for the User Model initialization. Extracted information is incorporated into the User Model.

Because the User Model is constantly and dynamically updated, applying the initialization process for each update is inefficient. Preferably, incremental learning techniques are used to update the User Model. Efficient incremental learning and updating techniques provide for incorporating new items into existing statistics, as long as sufficient statistics are recorded. Details about incremental learning can be found in P. Lee, *Bayesian Statistics*, Oxford University Press, 1989.

After a document stored in the recently accessed buffer is parsed, parsed portions are stored in candidate tables. For example, FIGS. 15A and 15B illustrate a user site candidate table and user word candidate table. The user site candidate table holds sites that are candidates to move into the user site distribution of FIG. 4B. The site candidate table stores the site name, i.e., the URL until the first backslash, except for special cases; the number of site accesses; and the time of last access. The user word candidate table holds the words or phrases that are candidates to move into the user informative word list of FIG. 4A. It contains a word or phrase ID, alternate spellings (or misspellings) of the word, an informative grade, and a time of last access.

Negative examples provide words, sites, and topics that can be used in several ways. The measure of any item obtained from the negative example may be reduced in the user distribution. For example, if the negative example is from a particular site that is in the user site distribution, then the probability or mutual information of that site is decreased. Alternatively, a list of informative negative items may be stored. The negative items are obtained from negative examples and are used to reduce the score of a document containing negative items.

Documents are added to the buffer during all user modes of interaction with the computer. Interaction modes include network searching, network navigation, network browsing, email reading, email writing, document writing, viewing "pushed" information, finding expert advice, and product

purchasing. Different types of information are stored in the buffer for different modes. In network searching, search queries are recorded and all search results added to the buffer, along with whether or not a link was followed and access time for viewed search results. In network browsing, the user browses among linked documents, and each document is added to the buffer, along with its interaction time. In email reading mode, each piece of email is considered to be a document and is added to the buffer. The type of interaction with the email item, such as deleting, storing, or forwarding, the sender of the email, and the recipient list are recorded. In email writing mode, each piece of written email is considered a document and added to the buffer. The recipient of the email is recorded. Documents written during document writing mode are added to the buffer. The user's access time with each piece of pushed information and type of interaction, such as saving or forwarding, are recorded. In finding expert advice mode, the user's interest in expert advice is recorded; interest may be measured by the interaction time with an email from an expert, a user's direct rating of the quality of information received, or other suitable measure.

During a product purchasing mode, a similar buffer is created for purchased products, as shown in FIG. 16. All purchased products are used to update the User Model. The user recently purchased products buffer records for each purchase the product ID, parent node in the product tree, purchase time, and purchase source. Purchased products are used to update the user product distribution and user product feature distribution.

If the user feels that the User Model is not an adequate representation of him or her, the user may submit user modification requests. For example, the user may request that specific web sites, topics, or phrases be added to or deleted from the User Model.

User Models for prototype users (hats) are also updated based on actions of similar users. Obviously, it is desirable for prototype User Models to reflect the current state of the representative interest. New web sites appear constantly, and even new informative words appear regularly. For example, technology-related words are introduced and widely adopted quite rapidly; the word list of the Technologist hat should be updated to reflect such changes.

Prototype User Models are updated using actions that are related to the prototype. Actions include documents, user reactions to documents, and product purchases. There are many ways to determine whether an action is relevant to the prototype user. A document that is a positive example for many users (i.e., a followed search result or bookmarked page) and also has a high probability of interest to the prototype user is added to the set of prototype user documents. Actions of users or clusters who are similar to the prototype user, as measured by the relative entropy between individual distributions (words, sites, etc.), are incorporated into the prototype User Model. Additions to the prototype User Model may be weighted by the relative entropy between the user performing the action and the prototype user. Actions of expert users who have a high degree of interest in topics also of interest to the prototype user are incorporated into the prototype User Model.

Note that users who are trying on hats are not able to change the prototype User Model. Their actions affect their own User Models, but not the prototype User Model. Updates to the prototype User Model are based only on actions of users who are not currently trying on hats.

Product models are also continually updated using incremental learning techniques. As described below, the present invention includes crawling network documents and evaluat-

ing each document against User Models. Crawled documents are also evaluated by product models. Documents that are relevant to a particular product, as determined by the computed probability $P(p|d)$, are used to update its product model. If a document is determined to be relevant, then each component of the product model is updated accordingly. In addition to the parsing and analysis performed for user documents, information extraction techniques are employed to derive feature values that are compared against feature values of the product model, and also incorporated into the feature value list as necessary. New products can be added to the product tree at any time, with characteristic product feature values extracted from all relevant documents. Relevant documents for updating product models include product releases, discussion group entries, product reviews, news articles, or any other type of document.

By employing dynamically updated product models, the present invention, in contrast with prior art systems, provides for deep analysis of all available product information to create a rich representation of products. The interest of a user in a product can therefore be determined even if the product has never been purchased before, or if the product has only been purchased by a very small number of users.

Applying the User Model to Unseen Documents

The User Model is applied to unseen documents to determine the probability that a document is of interest to the user, or the probability that a document is of interest to a user in a particular context. The basic functionality of this determination is then used in the various applications described in subsequent sections to provide personalized information and product services to the user.

The process of estimating user interest in a particular unseen document **120** is illustrated in FIG. 17. This process has the following three steps:

1. Preprocessing the document as for initialization (step **122**).
2. Calculating an individual score for the document for each element of the user representation (e.g., topic distribution, word list).
3. Non-linearly combining (**124**) individual scores into one score **126**, the probability that the user is interested in the unseen document, $P(u|d)$.

The second step varies for each individual score. From the parsed text, the words of the document **120** are intersected with the words or phrases in the user informative word list **128**. For every word or phrase in common, the stored mutual information between the two indicator variables I_w and I_u is summed to obtain the word score. Alternatively, the TFIDF associated with the word are averaged for every common word or phrase. The location score is given by the probability that the document site is of interest to the user, based on the user site distribution **130**.

The topic classifiers **132** are applied to document **120** to determine the probability that the document relates to a particular topic, $P(t|d)$. The user topic score is obtained by computing the relative entropy between the topic distribution $P(t|d)$ and the user topic distribution **134**, $P(t|u)$. After the document has been classified into topics, the topic expert models **136** are applied as described above to determine a score reflecting the interest of users that are experts in the particular topics of this document.

Similarly, the product models **138** are applied to document **120** to determine which products or product categories it describes, $P(p|d)$. From the document product distribution, the product score is obtained by computing the relative entropy between the document product distribution and user product distribution **140**, $P(p|u)$. For each product having a

nonzero value of $P(p|d)$, its feature values are given by the product model. The user's measures on each of these feature values, found in the user product feature distribution **141**, are averaged to obtain a product feature score for each relevant product. Product feature scores are then averaged to obtain an overall product feature score.

The cluster models **142** of clusters to which the user belongs are applied to the document to obtain $P(c(u)|d)$. This group model represents the average interests of all users in the cluster. Conceptually, the cluster model is obtained from the union of all the member users' documents and product purchases. Practically, the cluster model is computed from the User Models by averaging the different distributions of the individual User Models, and not from the documents or purchases themselves. Note that in a recursive way, all users have some impact (relative to their similarity to the user under discussion) on the user score, given that $P(c(u)|d)$ is estimated using $P(c(c(u))|d)$ as a knowledge source, and so on.

Finally, world knowledge (not shown) is an additional knowledge source that represents the interest of an average user in the document based only on a set of predefined factors. World knowledge factors include facts or knowledge about the document, such as links pointing to and from the document or metadata about the document, for example, its author, publisher, time of publication, age, or language. Also included may be the number of users who have accessed the document, saved it in a favorites list, or been previously interested in the document. World knowledge is represented as a probability between 0 and 1.

In step **124**, all individual scores are combined to obtain a composite user score **126** for document **120**. Step **124** may be performed by training a Multilayer Perceptron using jack-knifing and cross-validation techniques, as described in H. Bourlard and N. Morgan, *Connectionist Speech Recognition: A Hybrid Approach*, Kluwer Academic Publishers, 1994. It has been shown in J. Hertz et al., *Introduction to The Theory of Neural Computation*, Addison-Wesley, 1991, that a Multilayer Perceptron can be trained to estimate posterior probabilities.

The context of a user's interaction can be explicitly represented in calculating the user interest in a document. It is not feasible to update the user model after every newly viewed document or search, but the User Model can be updated effectively instantaneously by incorporating the context of user interactions. Context includes content and location of documents viewed during the current interaction session. For example, if the user visits ten consecutive sites pertaining to computer security, then when the User Model estimates the interest of the user in a document about computer security, it is higher than average. The probability of user interest in a document within the current context con is given by:

$$P(u|d, con) = \frac{P(u, con|d)}{P(con|d)}$$

In some applications, individual scores that are combined in step **124** are themselves useful. In particular, the probability that a user is interested in a given product can be used to suggest product purchases to a user. If a user has previously purchased a product, then the User Model contains a distribution on the product's features. If these features propagate far up the product tree, then they can be used to estimate the probability that the user is interested in a different type of product characterized by similar features. For example, if the user purchases a digital camera that is Windows compatible,

then the high probability of this compatibility feature value propagates up the tree to a higher node. Clearly, all computer-related purchases for this user should be Windows compatible. Every product that is a descendent of the node to which the value propagated can be rated based on its compatibility, and Windows-compatible products have a higher probability of being of interest to the user.

The long-term interest of a user in products, represented by $P(plu)$, is distinct from the user's immediate interest in a product p , represented as $P(uld, \text{product described}=p)$. The user's immediate interest is the value used to recommend products to a user. Note that $P(plu)$ does not incorporate the user's distribution on feature values. For example, consider the problem of evaluating a user's interest in a particular camera, the Nikon 320. The user has never read any documents describing the Nikon 320, and so $P(\text{Nikon 320}|u)=0$. However, the user's feature distribution for the Cameras node indicates high user interest in all of the feature values characterizing the Nikon 320.

When a given product is evaluated by the User Model, the following measures are combined to obtain $P(uld, \text{product described}=p)$: the probabilities of the product and its ancestor nodes from the user product distribution, $P(plu)$; an average of probabilities of each feature value from the user product feature distribution, $P(flu,p)$; a probability from the user's clusters' product distributions, $P(flc(u))$; and an average of probabilities of feature values from the cluster' product feature distributions, $P(flc(u),p)$. The overall product score is determined by non-linearly combining all measures. The cluster model is particularly useful if the user does not have a feature value distribution on products in which the user's interest is being estimated.

Applications

The basic function of estimating the probability that a user is interested in a document or product is exploited to provide different types of personalized services to the user. In each type of service, the user's response to the service provided is monitored to obtain positive and negative examples that are used to update the User Model. Example applications are detailed below. However, it is to be understood that all applications employing a trainable User Model as described above are within the scope of the present invention.

Personal Search

In this application, both the collection and filtering steps of searching are personalized. A set of documents of interest to the user is collected, and then used as part of the domain for subsequent searches. The collected documents may also be used as part of the user documents to update the User Model. The collection step, referred to as Personal Crawler, is illustrated schematically in FIG. **18**. A stack **170** is initialized with documents of high interest to the user, such as documents in the bookmarks file or documents specified by the user. If necessary, the stack documents may be selected by rating each document in the general document index according to the User Model. The term "stack" refers to a pushdown stack as described in detail in R. Sedgewick, *Algorithms in C++, Parts 1-4*, Addison-Wesley, 1998.

In step **172**, the crawler selects a document from the top of the stack to begin crawling. The document is parsed and analyzed (step **174**) to identify any links to other documents. If there are links to other documents, each linked document is scored using the User Model (**176**). If the linked document is of interest to the user (**178**), i.e., if $P(u|d)$ exceeds a threshold level, then it is added to the stack in step **180**, and the crawler continues crawling from the linked document (step **172**). If

the document is not of interest to the user, then the crawler selects the next document on the stack to continue crawling.

The subsequent searching step is illustrated in FIG. 19. In response to a query 190, a set of search results is located from the set containing all documents D and user documents obtained during personal crawling. The results are evaluated using the User Model (194) and sorted in order of user interest (196), so that the most interesting documents are listed first. The user reaction to each document in the search results is monitored. Monitored reactions include whether or not a document was viewed or ignored and the time spent viewing the document. Documents to which the user responds positively are parsed and analyzed (200) and then used to update the User Model (202) as described above.

The role of the User Model in filtering the search results in step 194 is based on Bayesian statistics and pattern classification theory. According to pattern classification theory, as detailed in R. Duda and P. Hart, *Pattern Classification and Scene Analysis*, Wiley, 1973, the optimal search result is the one with the highest posterior probability. That is, the optimal result is given by:

$$\text{Max}_D P(u|q, d),$$

where $P(u|q,d)$ is the posterior probability of the event that a document d is of interest to a user u having an information need q . This probability can be expressed as:

$$P(u|q, d) = \frac{P(q|d, u)P(u|d)}{P(q|d)}.$$

The term $P(u|d)$ represents the user interest in the document regardless of the current information need, and is calculated using the User Model. The term $P(q|d,u)$ represents the probability that a user u with an information need of d expresses it in the form of a query q . The term $P(q|d)$ represents the probability that an average user with an information need of d expresses it in the form of a query q . One possible implementation of the latter two terms uses the Hidden Markov Model, described in Christopher D. Manning and Hinrich Schütze, *Foundations of Statistical Natural Language Processing*, MIT Press, 1999.

Search results may also be filtered taking into account the context of user interactions, such as content of a recently viewed page or pages. When the context is included, the relevant equation is:

$$P(u|q, d, con) = \frac{P(q|d, u, con)P(u|d, con)}{P(q|d, con)},$$

where $P(u|d,con)$ is as described above.

The Personal Crawler is also used to collect and index documents for product models. Collected documents are parsed and analyzed to update product models, particularly the list of product feature values, which are extracted from collected documents using information extraction techniques.

In general, searches are performed to retrieve all documents from the set of indexed documents that match the search query. Alternatively, searches can be limited to prod-

uct-related documents, based on either the user's request, the particular search query, or the user's context. For example, a user is interested in purchasing a new bicycle. In one embodiment, the user selects a check-box or other graphical device to indicate that only product-related documents should be retrieved. When the box is not checked, a search query "bicycle" returns sites of bicycle clubs and newsletters. When the box is checked, only documents that have a nonzero product probability ($P(p|d)$) on specific products are returned. Such documents include product pages from web sites of bicycle manufacturers, product reviews, and discussion group entries evaluating specific bicycle models.

Alternatively, the search query itself is used to determine the type of pages to return. For example, a query "bicycle" again returns sites of bicycle clubs and newsletters. However, a query "cannondale bicycle" or "cannondale" returns only product-related pages for Cannondale bicycles. Alternatively, the user's context is used to determine the type of pages to return. If the last ten pages viewed by the user are product-related pages discussing Cannondale bicycles, then the query "bicycle" returns product-related pages for all brands of bicycles that are of interest to the user, as determined by the User Model. In all three possible embodiments, within the allowable subset of documents, the entire document is evaluated by the User Model to estimate the probability that the user is interested in the document.

Searches may also be performed for products directly, and not for product-related documents. Results are evaluated using only the user product distribution, user product feature distribution, and product and feature distributions of the user's clusters, as explained above. In general, product searches are performed only at the request of the user, for example by selecting a "product search" tab using a mouse or other input device. A user enters a product category and particular feature values, and a list of products that are estimated to be of high interest to the user is returned. The user is returned some form of list of most interesting products. The list may contain only the product name, and may include descriptions, links to relevant documents, images, or any other appropriate information.

Personal Browsing and Navigation

The present invention personalizes browsing and navigation in a variety of different ways. In the personal web sites application, web sites located on third party servers are written in a script language that enables dynamic tailoring of the site to the user interests. Parameters of the User Model are transferred to the site when a user requests a particular page, and only selected content or links are displayed to the user. In one embodiment, the site has different content possibilities, and each possibility is evaluated by the User Model. For example, the CNN home page includes several potential lead articles, and only the one that is most interesting to the user is displayed. In a second embodiment, links on a page are shown only if the page to which they link is of interest to the user. For example, following the lead article on the CNN home page are links to related articles, and only those of interest to the user are shown or highlighted. One single article has a variety of potential related articles; a story on the Microsoft trial, for example, has related articles exploring legal, technical, and financial ramifications, and only those meeting the user's information needs are displayed.

The personal links application is illustrated in FIG. 20. In this application, the hyperlinks in a document being viewed by the user are graphically altered, e.g., in their color, to indicate the degree of interest of the linked documents to the user. As a user views a document (step 210), the document is

parsed and analyzed (212) to locate hyperlinks to other documents. The linked documents are located in step 214 (but not shown to the user), and evaluated with the User Model (214) to estimate the user's interest in each of the linked documents. In step 216, the graphical representation of the linked documents is altered in accordance with the score computed with the User Model. For example, the links may be color coded, with red links being most interesting and blue links being least interesting, changed in size, with large links being most interesting, or changed in transparency, with uninteresting links being faded. If the user follows one of the interesting links (218), then the process is repeated for the newly viewed document (210).

The personal related pages application locates pages related to a viewed page. Upon the user's request (e.g., by clicking a button with a mouse pointer), the related pages are displayed. Related pages are selected from the set of user documents collected by the personal crawler. Implementation is similar to that of the personal search application, with the viewed page serving as the query. Thus the relevant equation becomes

$$P(u | \text{page}, d) = \frac{P(\text{page} | d, u)P(u | d)}{P(\text{page} | d)},$$

with $P(\text{page} | d, u)$ representing the probability that a user u with an information need of document d expresses it in the form of the viewed page page . $P(\text{page} | d)$ represents the probability that an average user with an information need of document d expresses it in the form of the viewed page page . These terms can be calculated using the Hidden Markov Model.

Alternatively, related pages or sites may be selected according to the cluster model of clusters to which the user belongs. The most likely site navigation from the viewed site, based on the behavior of the cluster members, is displayed to user upon request.

Related pages are particularly useful in satisfying product information needs. For example, if the user is viewing a product page of a specific printer on the manufacturer's web site, clicking the "related pages" button returns pages comparing this printer to other printers, relevant newsgroup discussions, or pages of comparable printers of different manufacturers. All returned related pages have been evaluated by the User Model to be of interest to the user.

Find the Experts

In this application, expert users are located who meet a particular information or product need of the user. Expert users are users whose User Model indicates a high degree of interest in the information need of the user. The information need is expressed as a document or product that the user identifies as representing his or her need. In this context, a document may be a full document, a document excerpt, including paragraphs, phrases, or words, the top result of a search based on a user query, or an email message requesting help with a particular subject. From the pool of potential experts, User Models are applied to the document or product, and users whose probability of interest in the document or product exceeds a threshold level are considered expert users.

The pool of experts is specified either by the user or in the system. For example, the pool may include all company employees or users who have previously agreed to help and advise other users. When users request expert advice about a particular product, the expert may be chosen from the product

manufacturer or from users who have previously purchased the product, or from users participating in discussion groups about the product.

A protocol for linking users and identified experts is determined. For example, the expert receives an email message requesting that he or she contact the user in need of assistance. Alternatively, all user needs are organized in a taxonomy of advice topics, and an expert searches for requests associated with his or her topic of expertise.

Personal News

This application, also known as personal pushed information, uses the personal crawler illustrated in FIG. 18. From all documents collected within a recent time period by the user's crawler or user's clusters' crawlers, the most interesting ones are chosen according to the User Model. Collection sources may also be documents obtained from news wires of actions of other users. Documents are sent to the user in any suitable manner. For example, users receive email messages containing URLs of interesting pages, or links are displayed on a personal web page that the user visits.

Personalization Assistant

Using the User Model, the Personalization Assistant can transform any services available on the web into personalized services, such as shopping assistants, chatting browsers, or matchmaking assistants.

Document Barometer

The document barometer, or Page-O-Meter, application, illustrated in FIG. 21, finds the average interest of a large group of users in a document. The barometer can be used by third parties, such as marketing or public relations groups, to analyze the interest of user groups in sets of documents, advertising, or sites, and then modify the documents or target advertising at particular user groups. The application can instead report a score for a single user's interest in a document, allowing the user to determine whether the system is properly evaluating his or her interest. If not, the user can make user modification requests for individual elements of the User Model. From individual and average scores, the application determines a specific user or users interested in the document.

Referring to FIG. 21, a document 220 is parsed and analyzed (222) and then evaluated according to a set of N User Models 224 and 226 through 228. N includes any number greater than or equal to one. The resulting scores from all User Models are combined and analyzed in step 230. In one embodiment, the analysis locates users having maximum interest in document 220, or interest above a threshold level, and returns a sorted list of interested users (232). Alternatively, an average score for document 220 is calculated and returned (234). The average score may be for all users or for users whose interest exceeds a threshold interest level. The range of interest levels among all users in the group may also be reported.

An analogous product barometer calculates user interest in a product. The product barometer computes a score for an individual user or group of users, or identifies users having an interest in a product that exceeds a threshold level. Third party organizations use the product barometer to target marketing efforts to users who are highly likely to be interested in particular products.

3D Map

FIG. 22 illustrates a three-dimensional (3D) map 240 of the present invention, in which rectangles represent documents and lines represent hyperlinks between documents. A user provides a set of hyperlinked documents, and each document

31

is scored according to the User Model. An image of 3D map **240** is returned to the user. 3D map **240** contains, for each document, a score reflecting the probability of interest of the user in the document.

Product Recommendations

A user's online shopping experience can be personalized by making use of the user's overall product score described above, $P(\text{uld, product described}=\text{p})$. Products that are of high interest to the user are suggested to him or her for purchase. When a user requests information for a specific product or purchases a product, related products are suggested (up-sell). Related product categories are predetermined by a human, but individual products within related categories are evaluated by the User Model before being suggested to the user. The related products are given to the user in a list that may contain images, hyperlinks to documents, or any other suitable information. For example, when a user purchases a server, a list of relevant backup tapes are suggested to him or her for purchase. Suggested products may have feature values that are known to be of interest to the user, or may have been purchased by other members of the user's cluster who also purchased the server. Related product suggestions may be made at any time, not only when a user purchases or requests information about a particular product. Suggested products may be related to any previously purchased products.

Similarly, competing or comparable products are suggested to the user (cross-sell). When the user browses pages of a particular product, or begins to purchase a product, products within the same product category are evaluated to estimate the user's interest in them. Products that are highly interesting to the user are recommended. The user might intend to purchase one product, but be shown products that are more useful or interesting to him or her.

It will be clear to one skilled in the art that the above embodiments may be altered in many ways without departing from the scope of the invention. Accordingly, the scope of the invention should be determined by the following claims and their legal equivalents.

What is claimed is:

1. A computer-implemented method for providing personalized information services to a user, the method comprising: transparently monitoring user interactions with data while the user is engaged in normal use of a browser program running on the computer; analyzing the monitored data to determine documents of interest to the user; estimating parameters of a user-specific learning machine based at least in part on the documents of interest to the user; receiving a search query from the user; retrieving a plurality of documents based on the search query; for each retrieved document of said plurality of retrieved documents: identifying properties of the retrieved document, and applying the identified properties of the retrieved document to the user-specific learning machine to estimate a probability that the retrieved document is of interest to the user; and using the estimated probabilities for the respective plurality of retrieved documents to present at least a portion of the retrieved documents to the user.
2. The method of claim 1, further comprising presenting to said user a list of said portion of the retrieved documents.

32

3. The method of claim 1, wherein transparently monitoring user interactions with data comprises monitoring user interactions with data during multiple different modes of user interaction with network data.

4. The method of claim 3, wherein the multiple different modes of user interaction comprise a plurality of modes selected from the group consisting of a network searching mode, a network navigation mode, and a network browsing mode.

5. The method of claim 1, further comprising analyzing the monitored data to determine documents not of interest to the user, and wherein estimating parameters of a user-specific learning machine further comprises estimating parameters of a user-specific learning machine based at least in part on the documents not of interest to the user.

6. The method of claim 1, wherein monitoring user interactions with data for a document comprises monitoring at least one type of data selected from the group consisting of information about the document, whether the user viewed the document, information about the user's interaction with the document, context information, the user's degree of interest in the document, time spent by the user viewing the document, whether the user followed at least one link contained in the document, and a number of links in the document followed by the user.

7. The method of claim 1, wherein said plurality of retrieved documents correspond to a respective plurality of products.

8. The method of claim 7, wherein using the estimated probabilities to present at least a portion of the retrieved documents to the user comprises presenting at least a portion of said products to the user.

9. The method of claim 1, wherein said search query pertains to a product of interest to the user, and wherein retrieving said plurality of documents based on the search query comprises retrieving a plurality of documents pertaining to a plurality of products related to the product of interest to the user.

10. The method of claim 9, wherein applying the identified properties of the retrieved document comprises applying the identified properties of the retrieved document pertaining to said related product to the user-specific learning machine to estimate a probability that the related product is of interest to the user.

11. The method of claim 10, wherein using the estimated probabilities for the respective plurality of retrieved documents comprises using the estimated probabilities for the respective plurality of retrieved documents pertaining to the related products to present at least a portion of the related products to the user.

12. The method of claim 1, further comprising estimating parameters of said user-specific learning machine based on a set of initial parameters identified at least in part on initial documents associated with said browser program.

13. The method of claim 12, wherein said initial documents are selected from the group of files consisting of favorites, bookmarks, cached files, temporary Internet files, and browsing history.

14. The method of claim 1, wherein identifying properties of the retrieved document comprises determining whether at least one of said documents of interest contains a link to said retrieved document.

15. The method of claim 1, wherein at least one of said properties of the retrieved document is based on intermediate documents linking from at least one of said documents of interest to said user towards said retrieved document.

33

16. The method of claim 15, wherein identifying properties of the retrieved document further comprises estimating a probability that at least one of said intermediate document linking from at least one of said documents of interest to said user towards said retrieved document are of interest to the user. 5

17. The method of claim 1, wherein identifying properties of the retrieved document further comprises estimating a probability that at least one intermediate document linking from at least one of said documents of interest to said user towards said retrieved document are of interest to the user. 10

18. The method of claim 1, wherein analyzing the monitored data to determine documents of interest to the user comprises analyzing said monitored data to obtain data associated with said monitored data selected from the group consisting of text, images, non-text media, and formatting. 15

19. The method of claim 18, wherein identifying properties of the retrieved document comprises analyzing said retrieved document to obtain data associated with the retrieved document said associated data selected from the group consisting of text, images, non-text media, and formatting. 20

20. The method of claim 19, wherein applying the identified properties of the retrieved document to the user-specific learning machine comprises comparing said data associated with said retrieved document with data in said user-specific learning machine having a type corresponding thereto. 25

21. The method of claim 1, wherein using the estimated probabilities for the respective plurality of retrieved documents to present at least a portion of the retrieved documents to the user comprises presenting to the user at least said portion of the retrieved documents based on the estimated probability that the retrieved document is of interest to the user and the relevance of the retrieved document to the search query. 30

22. The method of claim 1, wherein identifying properties of the retrieved document comprises identifying properties selected from the properties consisting of a topic associated with the retrieved document, at least one product feature extracted from the retrieved document, an author of the retrieved document, an age of the retrieved document, a list of documents linked to the retrieved document, a number of users who have accessed the retrieved document, and a number of users who have saved the retrieved document in a favorite document list. 35 40

23. A computer-implemented method for providing personalized information services to a user, the method comprising: 45

34

transparently monitoring user interactions with data while the user is engaged in normal use of a browser program running on the computer;

analyzing the monitored data to determine documents of interest to the user;

estimating parameters of a user-specific learning machine based at least in part on the documents of interest to the user;

collecting a plurality of documents of interest to a user; for each of said plurality of collected documents: identifying properties of the collected document, and applying the identified properties of the collected document to the user-specific learning machine to estimate a probability that the collected document is of interest to the user;

using the estimated probabilities for the respective plurality of collected documents to select at least a portion of the collected documents;

presenting said selected collected documents to said user.

24. The method of claim 23, wherein presenting said selected collected documents to said user comprises displaying said selected collected documents to said user on a personal web page associated with the user.

25. The method of claim 23, wherein said plurality of collected documents correspond to a respective plurality of products.

26. The method of claim 25, wherein using the estimated probabilities to present at least a portion of the retrieved documents to the user comprises presenting at least a portion of said products to the user.

27. The method of claim 24, wherein analyzing the monitored data to determine documents of interest to the user comprises analyzing said monitored data to obtain data associated with said monitored data selected from the group consisting of text, images, non-text media, and formatting.

28. The method of claim 27, wherein identifying properties of the collected document comprises analyzing said collected document to obtain data associated with the collected document said associated data selected from the group consisting of: text, images, non-text media, and formatting.

29. The method of claim 28, wherein applying the identified properties of the collected document to the user-specific learning machine comprises comparing said data associated with said collected document with data in said user-specific learning machine having a type corresponding thereto.

* * * * *

EXHIBIT 3

PUM v. Google

DISPUTED CLAIM CONSTRUCTION COMPARISON

Item #	Claim Term/Phrase	Claim(s)	P.U.M.'s Construction	Google's Construction
1	order of steps	'040 Patent: 1, 32 '276 Patent: 1, 23	No construction needed. If the Court is inclined to address the issue, then it should hold that the steps may be performed in a consecutive manner, in an overlapping manner, or a combination of the two, except as set forth below.	'040 Patent, 1 and 32: Steps (a), (b), and (c) must be performed in that order and before steps (e) and (f); step (d) must be performed before steps (e) and (f); and step (e) must be performed before step (f). '276 Patent, 1: steps (a), (b), and (c) in that order; step (d) before step (e); step (f) must be performed after steps (c) and (e); and step (g) must be performed after step (f). '276 Patent, 23: step (a), (b), (c), (d), (e), and (f) in that order
	antecedent basis terms			
	“User u”/ “the user” and “the user u”	'040 Patent: 1, 11, 21, 34 (depending on claim 32)	No construction necessary	“A user u” and “the user” / “the user u” refer to the same user.
	“user” / “the user”	'276 Patent: 1, 6, 21, 23	No construction necessary	“A user” and “the user” refer to the same user
	“user-specific data files” / “the user-specific data files”	'040 Patent: 1	No construction necessary	“user-specific data files” and “the user-specific data files” refer to the same files
	“a document d” / “the document”	'040 Patent: 1, 34 (depending on claim 32)	No construction necessary	“a document d” and “the document” refer to the same document.
	“a document” / “the document”	'276 Patent: 6	No construction necessary	“a document” and “the document” refer to the same document
	“a learning machine”/ “the learning machine”	'040 Patent: 1, 34 (depending on claim 32)	No construction necessary	“a learning machine” and “the learning machine” refer to the same learning machine.
	“a user-specific learning machine” / “the user-specific learning machine”	'276 Patent: 1, 23	No construction necessary	“a user-specific learning machine” and “the user-specific learning machine” refer to the same user-specific learning machine
	“a probability P(u d) that an unseen document d is of interest to the user u” / “the probability P(u d)” / “the	'040 Patent: 1, 34 (depending on claim 32)	No construction necessary	“a probability P(u d) that an unseen document d is of interest to the user u,” “the probability P(u d),” and “the estimated probability” refer to

Item #	Claim Term/Phrase	Claim(s)	P.U.M.'s Construction	Google's Construction
	estimated probability”			the same probability.
	“parameters of a learning machine” / “the parameters”	'040 Patent: 1,21, 34 (depending on claim 32)	No construction necessary	“parameters of a learning machine” and “the parameters” refer to the same parameters.
	“a user model” / “the user model”	'276 Patent: 1, 23	No construction necessary	“a user model” and “the user model” refer to the same user model.
	“a search query” / “the search query”	'276 Patent: 1, 23	No construction necessary	“a search query” and “the search query” refer to the same search query
2	“user” / “user [u]”	'040 patent: 1, 11, 32	“a person operating a computer as represented by a tag or identifier”	“person operating a computer”
3	“user-specific data files”	'040 Patent: 1, 32	“the monitored user interactions with data and a set of documents associated with the user”	“data files unique to the user”
	“monitored user interactions with the data”	'040 Patent: 1(b), 32	“the collected information about the user’s interactions with data”	“user interactions with data obtained from the monitoring step of 1[32](a)”
	“set of documents associated with the user”	'040 Patent: 1(b), 32	“a group or collection of text or other types of media associated with the user”	“group or collection of documents associated with the user”
4	“document”	<i>Passim</i>	“text or any type of media”	“an electronic file”
5	“estimating parameters of a learning machine”	040 Patent 1(c), 32(c)	“estimating values or weights of the variables of a learning machine”	“estimating a value or weight of each of the variables that are used by the learning machine to calculate a probability”
	“parameters of a learning machine”	'040 Patent: 1(c), 32(c)	<i>See construction for “estimating parameters of a learning machine”</i>	“variables, having a value or weight, used by the learning machine to calculate a probability”
	“estimating parameters of a user-specific learning machine”	'276 Patent: 1, 5, 23	“estimating values or weights of the variables of a user-specific learning machine”	“estimating a value or weight of each of the variables that are used by the learning machine to calculate a probability”
	“parameters of a user-specific learning machine”	'276 Patent: 1, 5, 23	<i>See construction from “estimating parameters of a user-specific learning machine”</i>	“variables, having a value or weight, used by the user specific learning machine to calculate a probability”
6	“learning machine”	<i>Passim</i>	“a model and/or mathematical function that is used to make a prediction or intelligent decision that attempts to improve performance in part by altering the values/weights given to its variables depending upon	“program that contains parameters used to calculate a probability, and where the predictive ability improves over time with the addition of new data.”

Item #	Claim Term/Phrase	Claim(s)	P.U.M.'s Construction	Google's Construction
			past observations or experiences”	
	“user-specific learning machine”	’276 Patent	“a model and/or mathematical function that is used to make a prediction or intelligent decision that attempts to improve performance in part by altering the values/weights given to its variables depending upon past observations or experiences specific to the user”	“learning machine unique to the user”
	“User Model specific to the user”	’040 Patent: 1(c), [21], 32	“an implementation of a learning machine updated in part from data specific to the user”	“model unique to the user, that is created and updated by the learning machine and stored in a data structure”
7	“estimating a probability P(u d) that an unseen document d is of interest to the user u”	’040 Patent: 1e, 32e	“approximating or roughly calculating the degree of belief or likelihood that an unseen document d is of interest to the user u given the information that is known about the unseen document”	“calculating the percentage chance that an unseen document d is of interest to the user u given the information that is known about the unseen document”
	“estimating a posterior probability P(u d,q) that the document d is of interest to the user u, given a query q submitted by the user”	’040 Patent: 11	“approximating or roughly calculating the degree of belief or likelihood that a document d is of interest to the user u given the information that is known about the document, and given a query q”	“calculating the percentage chance of the user u being interested, taking into account what is previously known about that user's interests in general, given new knowledge of the document d the user is considering and a search query q submitted by the user”
8	“unseen document”	’040 Patent: 1; 32	“document not previously seen by the user”	“document not previously seen by any user”
9	“present” or “presenting”	’276 Patent: 1, 21	“to provide or make available”	“display[ing]”
10	“user interest information derived from the User Model”	’040 Patent: 21, 52	“interests or other information inferred from the User Model”	Indefinite
	“documents of interest to the user” / “documents [that are] not of interest to the user”	’276 patent: 1, 5, 14, 23	“text or media for which the user has a positive response” / “text or media for which the user has a negative response or has ignored”	Indefinite

EXHIBIT 4

Searching the Web: New Domains for Inquiry

Bertram Bruce

University of Illinois at Urbana-Champaign
United States

This column is reprinted from the Technology department of the [*Journal of Adolescent & Adult Literacy*](#) (JAAL). It contains the following sections:

- [Author's Message](#)
 - [Issue of the Month](#)
 - [Data View](#)
 - [Glossary](#)
 - [References](#)
-

Author's Message

Many teachers today recognize the importance of online data sources for all kinds of research and writing projects. Some now permit students to include online sources in their work, and others go so far as to require the use of online sources.

There is a cornucopia of resources online. Reference tools include encyclopedias, dictionaries, and collections of quotes; libraries of poetry, short stories, images, and music; critical studies and research articles on every conceivable topic; information about authors and historical figures; government and public policy data; current events; and much, much more. Most teachers quickly see the problems that arise from such bounty. Issues of plagiarism, pornography, commercialism, and simple time wasting soon rear up regardless of the topic. When the cornucopia spills out 100,000 Web sites of dubious quality and relevance, it seems much less bounteous.

This month's Technology column addresses why it is important to think more critically about Web searching. Questions of quantity become important. As the Web grows rapidly, unpredictably, unevenly, and without the familiar monitors provided by textbook companies or district curriculum guides, how should we think about its use? For a start, how do the size of the Web and the quality of material on it affect searching? Given these issues, what are some good approaches to search the Web effectively? What tools are available and how can they be used?

These questions point to even more fundamental issues. Perhaps we need to move from a conception of searching the Web to find a piece of information to one in which a search is embedded in how we think. This leads to perhaps the most important question: How can searching become not only "looking up," but truly productive inquiring?

[Back to top](#)

[Back to listing of all JAAL Technology columns available in Reading Online](#)

Issue of the Month: Searching Is the Journey, Not Just the Arrival

When students search the Web, it often seems that the problems are greater than the rewards. We seek ways to control those searches to avoid objectionable materials, plagiarism, or aimlessness, but in the process we may miss what is most valuable about the Web. Let us consider some questions about searches, which may help resolve this quandary.

Why is it important to think critically about Web searching?

We tend to think first of Web searching as a simple process of looking up some item of information. For certain purposes, that conception is quite appropriate. For example, if I want to find general information about William Shakespeare, I type *Shakespeare* into my [search engine](#). I get 660,000 Web pages back, but among the top 10 is the [Folger Shakespeare Library](#) in Washington, D.C. (administered by the Trustees of Amherst College), which has all sorts of interesting information, including lesson plans for teaching Shakespeare.

But, suppose I want to enter into the critical debates about Shakespearean authorship. Among the top 10 is the home page of the [Shakespeare Oxford Society](#). This group claims to be

the second oldest continuously operating organization...involved in the...Shakespeare authorship debate. The purpose of the Society is to document and establish Edward de Vere, 17th Earl of Oxford (1550-1604), as the universally recognized author of the works of William Shakespeare.

I am inclined to believe their claim and am intrigued to examine their arguments. As I explore their Web site I am impressed by the care shown with their presentation, the detail of their documents, the source citations, and the opportunities for feedback. But, as a novice in this area I cannot be certain that this is the most credible starting point for my inquiry. Is this enterprise considered to be a fringe group? Are there more credible sources, perhaps even espousing the same argument?

Despite these concerns, I am relieved in some ways. Although the organization offers books and videos for sale through the site, these commercial aspects appear supportive of the generally academic mission. I don't see here troubling signs of racism or pornography that permeate the Web. My usual worry that the site may be superseded by a more recent one is allayed by the fact that their latest update is the date of my visit. Thus, although I do not know the authors or much about their domain of study, I find the site to be worth further investigation. If I can believe what I read there, I've found a timely resource with all sorts of useful information and links for further study.

What I have discovered here is a potentially useful source, but although I have spent some time examining it I still have doubts about how to interpret what I read there. When I return to the list of 660,000 documents that the search engine provided, I feel a bit overwhelmed. Will I have to spend this much time on every document and still not know what to make of it all? Will my students cope with this any better than I do?

I have discovered something else. For certain kinds of queries, my search is far from a simple "look-up." Instead, it appears to be part of the general process of inquiry, which is tentative and fallible. There is no absolute starting point nor any sure way to reach the end, assuming such a point exists. I need to muster all my resources for critical thinking to navigate the Web, but I may reap enormous benefits in the process.

How does the size of the Web affect searching?

The enormous size of the Web (see [Data View](#)) is a mixed blessing. Hundreds of millions of pages hold forth the promise of having the text or images we seek, but the sheer volume of material gets in its own way. I recently searched for the U.S. Department of Commerce's report "Falling Through the Net," which is about the racial and income inequities in access to new information and communication technologies. My search engine offered up articles about World Cup soccer and the

performance of an accomplished goalie defending her team's net. At other times I have found obsolete versions of material that exists elsewhere on the Web, but is unknown to the search engines. Very often I find it difficult to get past the many commercial sites that have engineered their Web pages to appear first no matter how I specify a search query.

Given the number of Web pages, it is surprising that one can find anything at all, much less do so in a matter of seconds. Improved search engines make that possible, especially when the user understands how the search engines work and puts some effort into selecting a good set of keywords.

How does the quality of material on the Web affect searching?

I heard a teacher say recently that she discourages students from using the Web for research because the quality of material there is so poor. Although I would not abandon the Web because of its negative features, I can certainly sympathize. In fact, I can imagine that she might provide a list like the one below to support her point.

- *Hate sites, pornography, violence, criminal activity, et cetera*
It is unfortunately the case that one cannot imagine any dark corner of human activity that is not now represented on the Web. Sites promoting substance abuse, suicide, bomb making, and racial hatred interleave with children's artwork, poetry, music, and images from the Hubble telescope.
- *Commercialism*
Too many Web sites are created to sell something, not to provide valid and useful information. A recent estimate is that 83% of sites have primarily commercial content ([Guernsey, 1999; online document](#), registration required for access). Even in cases where the information is useful, the commercial assault is something to be avoided in schools and other learning environments.
- *Incompleteness*
The Web pretends to a universality that it cannot support. It represents human knowledge, cultures, and values very unevenly, yet the hypertext medium suggests that everything is really there and equitably represented somehow.
- *Authority*
The beauty of the Web is that anyone can make a Web site, for less than the cost of publishing a pamphlet. But a consequence of this is that there is no resort to any kind of recognized textual authority and no board of editors (as for a respected encyclopedia) who invite authors and vet articles for publication.
- *Relevance*
There is so much material on the Web that the irrelevant far outweighs the relevant for any search.
- *Timeliness*
The [CNN Interactive Web site](#) is updated every few minutes. Other sites are created, posted on the Web, and never changed. Some sites, but not all, indicate when they were last updated, but it is usually difficult to determine whether the page you are viewing is the most recent in a series or just the one you happened upon. For Shakespeare, the timeliness issue may not be severe, but for many domains it is critical, yet unanswerable.
- *Plagiarism*
When students (or anyone, for that matter) do find relevant information on the Web, it is all too easy to copy without attribution.

What tools are available for searching and how can they be used?

Every technology arises out of the problems of previous technologies. This is a cycle we see operating with Web searches: The Web solves the problem of managing diverse, distributed sets of documents. That solution in turn makes it possible to post documents easily for all to read. This leads to a profusion of documents, many of which are poorly written and irrelevant for particular purposes. Search engines and search directories arise to solve the problem of managing the enormous quantity of material. Document designers then manipulate the pages that the search engines see so that their documents rank highest. [Filters](#) are developed to screen out unwanted material. Documents are designed to defeat the filters, and so on. A sampling of these tools are described in the [Glossary](#).

What are some good approaches to searching effectively?

Much work is now underway to build better search engines, search directories, filters, jump sites, portals, and other technologies to enable more productive use of the Web. But what can an individual do to improve the experience of using the Web?

There are Web sites (of course!) devoted to this question. For example, Terry Gray has a useful review of some of the top search engines at daphne.palomar.edu/TGSEARCH/ and provides search tips specific to each engine. The Community Learning Network has information about many search engines and subject directories and a good set of FAQs about searching at www.cln.org/searching_home.html. Instead of going into great detail, I'll highlight three basic principles: (a) Understand how the Web and searches work, (b) select appropriate tools, and (c) use those tools effectively.

On the first point, it must be said that no one fully understands the Web, and even if a few did they would find their knowledge quickly dated as the technology and Web content evolved. Nevertheless, it helps to know some basic facts about how the Web functions and how search tools can help navigation.

For example, search engines do not go out and look at every Web page to answer a query. Many pages are hidden from the search engines behind organizational [firewalls](#). Moreover, it would take far too long to examine every page as each query arises. Instead, the search engine builds a [search index](#) that enables fairly rapid searches. A consequence of this is that the user is not searching the Web, but the index, and is thus dependent on the quality of the index, its organizational scheme, and how recently it has been updated. Among other things, that means recent additions to the Web may not appear as the result of a search. A recent study ([Lawrence & Giles, 1999; online document](#)) found that the best of the search engines finds only 16% of the relevant Web pages, not counting those behind firewalls. These issues need to be understood when interpreting the results of a search.

The second point is that the choice of search engine or [search directory](#) is a major factor in how effective a search may be. For example, to find information about a book it may be more effective to search the database of an online bookseller than to search the entire Web. But, if the book is out of print it will not help to search the site of a bookseller offering only current titles.

Sites such as SearchIQ provide some information about the relative performance of different search tools, but there is no substitute for trying out different tools with the types of questions under investigation and then looking critically at the types of search results produced. It is also important to understand how a particular tool works and what assumptions it makes. A tool that aims to bring up frequently accessed sites may be appropriate if you plan to shop online and want to find popular commercial sites, but it is less appropriate if you want novel perspectives on understanding some issue of international relations.

Many people recommend [metasearch engines](#) such as Cyber411 at www.cyber411.com, which combines the result of 16 search engines. But, the larger number of hits may not offset the extra time that each search requires and the redundancy. This is particularly so because often if the desired sites do not appear in the first 10 items they might as well not appear at all.

The third point is to develop means for using these tools effectively. Each search engine has its own syntax for specifying Boolean expressions. Usually, a phrase in quotes means to find that phrase exactly as written. Thus, typing "*best search engine*" to [AltaVista](#) yields nearly 8,000 sites containing that phrase in quotes. Typing the three words *best search engine* without the quotes yields the same result, but typing *search engine best* produces 4.4 million Web pages, the intersection of the 1.5 million Web pages containing the term *search engine* and the 17.5 million containing the word *best*.

It is difficult to lay out general rules for doing searches because the approach depends on the problem being investigated. Perhaps one good general rule is this: If a search produces many irrelevant documents it is important to understand why that happened and not simply to decry the bloated Web world.

How can searching become not only looking up, but truly productive inquiring?

There are two problems with conceiving of Web searches as simply the looking up of information. The first is that we are often frustrated. The answers may be out there, but if we search inappropriately we get useless data. Most interesting questions require some effort ahead of time to be formulated well. It is worth giving a try to sites such as AskJeeves at www.askjeeves.com, but more often you will need to rethink the question in order to find the answers you seek.

The second problem is that the view of Web searching as simply finding information limits the key to its importance for education or other life activities. The joy and true value of the Web lie in the way it can open up our questions. We ask one thing, but the Web leads us to ask more questions and to become aware of how much we do not know. This suggests an alternative to the common practice of asking students to cite one library source and one online source for an essay. We could turn the Web's unruliness into a virtue. Instead, we might say, "Use the Web to find the answer to such-and-such question. Now, report on three things you learned that you had never imagined before you did that search."

The Web search engines are very important and useful resources, and they are playing a major role in the information age. However, they currently lack comprehensiveness and timeliness. The current state of search engines can be compared to a phone book that is updated irregularly, and has most of the pages ripped out ([Lawrence & Giles, 1999; online document](#)).

[Back to top](#)

[Back to listing of all JAAL Technology columns available in Reading Online](#)

Data View: How Big Is the Web?

Like many simple questions, this one turns out to be more complicated than it might at first appear. A good activity for students would be for them to define what they mean by the terms *big* and *Web* and then to search the Web for data or analyses to help them answer that question. Different approaches could lead to varying results, which in turn might call for justification of their strategies and critical thinking (see [Murphy, 1998; online document](#), registration required for access). There are a number of things we might count.

- *Users*

The Internet surveys by the [Nua company](#) of Dublin, Ireland, estimate that 179 million people

were Internet users as of June 1999. This is about 3% of the earth's population. More than half of those users are in the U.S. and Canada.

- [Hosts](#)

The Domain Survey attempts to discover every host on the Internet by doing a complete search of the [Domain Name](#) system to find the name assigned to every possible Internet Provider (IP) address. It is sponsored by the [Internet Software Consortium](#). They estimated in January 1999 that there were more than 43 million hosts on the Internet. The number of hosts had nearly doubled every year.

- [Web pages](#)

A Web page can be anything from a few words to a site with video, interactive software, music files, or extensive text. Thus, when we say that there are so many pages on the Web, it is not quite the same as saying so many pages in a book. Still, if we knew how many Web pages there were we would have some idea of the size of the Web, at least relative to what it has been.

People have developed a variety of ways to gauge the Web's size in terms of pages. By comparing the pages returned by various search engines, [Lawrence and Giles \(1999; online document\)](#) derived [800 million pages](#) as a lower bound for the size of the (publicly accessible) Web as of February 1999 (see also [Guernsey, 1999; online document](#), registration required for access). This means the Web contains at least 800 million pages, probably somewhat more. They estimated further that these pages contain 6 [terabytes](#) of data versus the 20 terabytes of the entire U.S. Library of Congress. By the time of this issue, that lower bound on the size of the Web should have increased to well over a billion pages.

You can easily do an experiment yourself to get a rough measure of the Web's size. Go to [AltaVista](#) and search for the word *the*. You won't get any hits, because the search algorithm ignores common function words, but you will see the number of pages that the algorithm ignores. In July 1999, I found about 2 billion pages by this method. This is a rough measure because it excludes the pages that AltaVista doesn't know about, counts some pages with invalid links, and double-counts others. Nevertheless, it is a reasonable estimate not too far off from some more complex approaches.

- [Hyperlinks](#)

[Members of the Clever Project \(1999; online document\)](#) estimate that roughly a million Web pages are added every day. This is one rationale for their effort to develop algorithms for more sophisticated searching. These algorithms are based on studies of the interconnectivity of the Web, on how Web pages have annotated connections to other pages. They estimate that there are more than a billion hyperlinks in the Web today, about one per page. If we take the higher estimate of Web pages given above, then the number of hyperlinks must be considerably larger.

Win Treese has an Internet index newsletter at <http://new-website.openmarket.com/intindex/index.cfm> that regularly reports interesting items about the size and growth of the Web in the manner of *Harper's Index*. You can visit the site or become a subscriber to the index.

Comprehensive information about search engines, specialized search engines, and metasearch engines, as well as general information about searching and tips for searching are available at [SearchIQ](#). This site provides independent reviews and rankings to inform the selection of search tools. Their reviews employ criteria such as overall relevancy of listings and organization by relevancy, ability to find sites for both broad and specific topics, comprehensiveness, lack of redundancy, logical grouping of listings, and speed.

Although the "IQ scores" that SearchIQ assigns are a novel feature of the site, I recommend looking

more at their descriptions of the search engines and how they work. A low-scoring engine could easily be better for some purposes than the top-ranked engine.

[Back to top](#)

[Back to listing of all JAAL Technology columns available in Reading Online](#)

Glossary

Authority -- a Web site that is linked from many other pages (see [Hub](#)).

Boolean expression -- an expression that evaluates to *true* or *false*; for example, used in a Web search, the expression *travel and France* is true for every Web page that contains both *travel* and *France*. Expressions that contain logical operators such as *and*, *or*, and *not* are Boolean, but all Web searches implicitly involve Boolean expressions.

[Back](#)

Case sensitive -- the property of paying attention to upper- and lower-case letters; each search engine has its own policy about this (e.g., is *White House* the same as *white house*?).

Domain name -- a name that identifies an IP address(es). For example, the domain name [www.ed.gov](#) represents a numerical address signifying a location in cyberspace. A domain name is the first part of the URL used to identify a Web page.

[Back](#)

Filter -- a program that takes a list of documents and removes those that meet certain prespecified criteria; family filters are used to remove objectionable Web material, other filters are used to focus a search to retrieve the most relevant items, and any filter will occasionally let through unwanted items and screen out desirable ones.

[Back](#)

Firewall -- a system that creates a partition between a private network and the larger Internet; it may restrict access both to and from the Internet.

[Back](#)

Host -- a term used to refer to any single machine on the Internet, but a single machine can act like multiple systems, each with its own domain name and IP address, and so the definition now typically includes virtual hosts as well.

[Back](#)

Hub -- a Web site with many links to other sites (see [Authority](#)).

Metasearch engine -- a computer program, such as [Dogpile](#), that collects the results from several search engines at once. This is especially valuable because no search engine indexes more than one sixth of the Web.

[Back](#)

Ranking function -- a means used by a search engine to order documents found in a search in terms of potential relevance, quality, or other criteria.

Search directory -- a database that organizes documents according to categories and, usually, subcategories; it provides an alternative to general searching for finding particular items.

[Back](#)

Search engine -- a computer program that returns a list of the documents that satisfy a [Boolean search expression](#); it's usually used to refer to programs that search for Web documents.

[Back](#)

Search index -- a large database of document locations based on the words contained in each document; the index facilitates efficient, meaningful searches and is created by a program within the search engine.

[Back](#)

Specialty search engine -- a search engine that searches a limited database of documents, such as the telephone white pages; such an engine can be made more efficient for limited purposes and is more likely to return only the sorts of data that a user would want.

Spiders (search robots) -- a computer program sent out by a search engine to find as many documents on the Web as it can.

Terabyte -- a trillion bytes of information, enough to represent a trillion characters; about 100 fairly large personal computer hard drives would be needed to hold this much information.

[Back](#)

Webopedia -- a good online glossary of computer terms.

[Back to top](#)

[Back to listing of all JAAL Technology columns available in Reading Online](#)

References

Guernsey, L. (1999, July 8). Seek -- But on the Web, you might not find. *New York Times on the Web*. Available: www.nytimes.com

[Back](#)

[Back \(2nd citation\)](#)

Lawrence, S., & Giles, C.L. (1999). How big is the Web? How much of the Web do the search engines index? How up to date are the search engines? Available:

<http://www.neci.nec.com/~lawrence/websize.html>

[Back](#)

[Back \(2nd citation\)](#)

[Back \(3rd citation\)](#)

Members of the Clever Project. (1999). Hypersearching the Web. *Scientific American*, 280(6), 54-60. Available: www.sciam.com/1999/0699issue/0699raghavan.html

[Back](#)

Murphy, J. (1998, July 5). It's not the size that counts, but how you measure it. *New York Times on the Web*. Available: www.nytimes.com

[Back](#)

[Back to top](#)

[Back to listing of all JAAL Technology columns available in Reading Online](#)

Reading Online, www.readingonline.org

Posted November 1999

Published simultaneously in the *Journal of Adolescent & Adult Literacy*

© 1999-2000 International Reading Association, Inc. ISSN 1096-1232

EXHIBIT 5

Manual of PATENT EXAMINING PROCEDURE

Original Eighth Edition, August 2001

Latest Revision July 2010



U.S. DEPARTMENT OF COMMERCE
United States Patent and Trademark Office

The U.S. Patent and Trademark Office does not handle the sale of the Manual, distribution of notices and revisions, or change of address of those on the subscription list. Correspondence relating to existing subscriptions should be sent to the Superintendent of Documents at the following address:

Superintendent of Documents
Mail List Section, Washington, DC 20402

Telephone: 202-512-2267

Inquiries relating to purchasing the Manual should be directed to:

Superintendent of Documents
United States Government Printing Office
Washington, DC 20402

Telephone: 202-512-1800

Orders for reproduced copies of individual replacement pages or of previous revisions of the Manual should be sent to the following address:

Mail Stop Document Services
Director of the U.S. Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

Telephone: 1-800-972-6382 or 571-272-3150

Previous editions and revisions of the Manual are available on microfilm in the Patent Search Room.
The Manual is available on CD-ROM and on diskette from:

U.S. Patent and Trademark Office
Office of Electronic Information Products
MDW 4C18, P.O. Box 1450
Alexandria, VA 22313-1450

Telephone: 571-272-5600

Employees of the U.S. Patent and Trademark Office should direct their requests for the Manual, replacement pages, notices, and revisions to the Office of Patent Training.

Telephone: 571-272-7222

Pursuant to the Patent and Trademark Office Efficiency Act (PTOEA) (Pub. L. 106-113, 113 Stat. 1501A-572), the head of the United States Patent and Trademark Office (USPTO) is the "Under Secretary of Commerce for Intellectual Property and Director of the United States Patent and Trademark Office." The Director is assisted by the "Deputy Under Secretary of Commerce for Intellectual Property and Deputy Director of the United States Patent and Trademark Office." The patent operations of the USPTO are now headed by the "Commissioner for Patents." The trademark operations of the USPTO are now headed by the "Commissioner for Trademarks." Under section 4741(b) of the PTOEA, any reference to the Commissioner of Patents and Trademarks, the Assistant Commissioner for Patents, or the Assistant Commissioner for Trademarks is deemed to refer to the Director, the Commissioner for Patents, or the Commissioner for Trademarks, respectively. See "Reestablishment of the Patent and Trademark Office as the United States Patent and Trademark Office" published in the *Federal Register* at 65 FR 17858 (Apr. 5, 2000), and in the *Official Gazette of the United States Patent and Trademark Office* at 1234 O.G. 41 (May 9, 2000).

Additions to the text of the Manual are indicated by arrows (><) inserted in the text. Deletions are indicated by a single asterisk (*) where a single word was deleted and by two asterisks (**) where more than one word was deleted. The use of three or five asterisks in the body of the laws, rules, treaties, and administrative instructions indicates a portion of the law, rule, treaty, or administrative instruction which was not reproduced.

First Edition, November 1949
Second Edition, November 1953
Third Edition, November 1961
Fourth Edition, June 1979
Fifth Edition, August 1983
Sixth Edition, January 1995
Seventh Edition, July 1998
Eighth Edition, August 2001
Revision 1, February 2003
Revision 2, May 2004
Revision 3, August 2005
Revision 4, October 2005
Revision 5, August 2006
Revision 6, September 2007
Revision 7, July 2008
Revision 8, July 2010

made. The examiner should analyze whether the metes and bounds of the claim are clearly set forth. Examples of claim language which have been held to be indefinite because the intended scope of the claim was unclear are:

(A) “R is halogen, for example, chlorine”;

(B) “material such as rock wool or asbestos” *Ex parte Hall*, 83 USPQ 38 (Bd. App. 1949);

(C) “lighter hydrocarbons, such, for example, as the vapors or gas produced” *Ex parte Hasche*, 86 USPQ 481 (Bd. App. 1949); and

(D) “normal operating conditions such as while in the container of a proportioner” *Ex parte Steigerwald*, 131 USPQ 74 (Bd. App. 1961).

>The above examples of claim language which have been held to be indefinite are fact specific and should not be applied as *per se* rules. See MPEP § 2173.02 for guidance regarding when it is appropriate to make a rejection under 35 U.S.C. 112, second paragraph.<

2173.05(e) Lack of Antecedent Basis [R-5]

A claim is indefinite when it contains words or phrases whose meaning is unclear. The lack of clarity could arise where a claim refers to “said lever” or “the lever,” where the claim contains no earlier recitation or limitation of a lever and where it would be unclear as to what element the limitation was making reference. Similarly, if two different levers are recited earlier in the claim, the recitation of “said lever” in the same or subsequent claim would be unclear where it is uncertain which of the two levers was intended. A claim which refers to “said aluminum lever,” but recites only “a lever” earlier in the claim, is indefinite because it is uncertain as to the lever to which reference is made. Obviously, however, the failure to provide explicit antecedent basis for terms does not always render a claim indefinite. If the scope of a claim would be reasonably ascertainable by those skilled in the art, then the claim is not indefinite. >*Energizer Holdings Inc. v. Int’l Trade Comm’n*, 435 F.3d 1366, 77 USPQ2d 1625 (Fed. Cir. 2006)(holding that “anode gel” provided by implication the antecedent basis for “zinc anode”);< *Ex parte Porter*, 25 USPQ2d 1144, 1145 (Bd. Pat. App. & Inter. 1992) (“controlled stream of fluid” provided reasonable antecedent basis for “the controlled fluid”). Inherent

components of elements recited have antecedent basis in the recitation of the components themselves. For example, the limitation “the outer surface of said sphere” would not require an antecedent recitation that the sphere has an outer surface. See *Bose Corp. v. JBL, Inc.*, 274 F.3d 1354, 1359, 61 USPQ2d 1216, 1218-19 (Fed. Cir 2001) (holding that recitation of “an ellipse” provided antecedent basis for “an ellipse having a major diameter” because “[t]here can be no dispute that mathematically an inherent characteristic of an ellipse is a major diameter”).

EXAMINER SHOULD SUGGEST CORRECTIONS TO ANTECEDENT PROBLEMS

Antecedent problems in the claims are typically drafting oversights that are easily corrected once they are brought to the attention of applicant. The examiner’s task of making sure the claim language complies with the requirements of the statute should be carried out in a positive and constructive way, so that minor problems can be identified and easily corrected, and so that the major effort is expended on more substantive issues. However, even though indefiniteness in claim language is of semantic origin, it is not rendered unobjectionable simply because it could have been corrected. *In re Hammack*, 427 F.2d 1384 n.5, 166 USPQ 209 n.5 (CCPA 1970).

A CLAIM TERM WHICH HAS NO ANTECEDENT BASIS IN THE DISCLOSURE IS NOT NECESSARILY INDEFINITE

The mere fact that a term or phrase used in the claim has no antecedent basis in the specification disclosure does not mean, necessarily, that the term or phrase is indefinite. There is no requirement that the words in the claim must match those used in the specification disclosure. Applicants are given a great deal of latitude in how they choose to define their invention so long as the terms and phrases used define the invention with a reasonable degree of clarity and precision.

A CLAIM IS NOT PER SE INDEFINITE IF THE BODY OF THE CLAIM RECITES ADDITIONAL ELEMENTS WHICH DO NOT APPEAR IN THE PREAMBLE

The mere fact that the body of a claim recites additional elements which do not appear in the claim’s

preamble does not render the claim indefinite under 35 U.S.C. 112, second paragraph. See *In re Larsen*, No. 01-1092 (Fed. Cir. May 9, 2001) (unpublished) (The preamble of the *Larsen* claim recited only a hanger and a loop but the body of the claim positively recited a linear member. The examiner rejected the claim under 35 U.S.C. 112, second paragraph, because the omission from the claim's preamble of a critical element (i.e., a linear member) renders that claim indefinite. The court reversed the examiner's rejection and stated that the totality of all the limitations of the claim and their interaction with each other must be considered to ascertain the inventor's contribution to the art. Upon review of the claim in its entirety, the court concluded that the claim at issue apprises one of ordinary skill in the art of its scope and, therefore, serves the notice function required by 35 U.S.C. 112, paragraph 2.).

2173.05(f) Reference to Limitations in Another Claim

A claim which makes reference to a preceding claim to define a limitation is an acceptable claim construction which should not necessarily be rejected as improper or confusing under 35 U.S.C. 112, second paragraph. For example, claims which read: "The product produced by the method of claim 1." or "A method of producing ethanol comprising contacting amylose with the culture of claim 1 under the following conditions" are not indefinite under 35 U.S.C. 112, second paragraph, merely because of the reference to another claim. See also *Ex parte Porter*, 25 USPQ2d 1144 (Bd. Pat. App. & Inter. 1992) where reference to "the nozzle of claim 7" in a method claim was held to comply with 35 U.S.C. 112, second paragraph. However, where the format of making reference to limitations recited in another claim results in confusion, then a rejection would be proper under 35 U.S.C. 112, second paragraph.

2173.05(g) Functional Limitations [R-3]

A functional limitation is an attempt to define something by what it does, rather than by what it is (e.g., as evidenced by its specific structure or specific ingredients). There is nothing inherently wrong with defining some part of an invention in functional terms. Functional language does not, in and of itself,

render a claim improper. *In re Swinehart*, 439 F.2d 210, 169 USPQ 226 (CCPA 1971).

A functional limitation must be evaluated and considered, just like any other limitation of the claim, for what it fairly conveys to a person of ordinary skill in the pertinent art in the context in which it is used. A functional limitation is often used in association with an element, ingredient, or step of a process to define a particular capability or purpose that is served by the recited element, ingredient or step. >In *Innova/Pure Water Inc. v. Safari Water Filtration Sys. Inc.*, 381 F.3d 1111, 1117-20, 72 USPQ2d 1001, 1006-08 (Fed. Cir. 2004), the court noted that the claim term "operatively connected" is "a general descriptive claim term frequently used in patent drafting to reflect a functional relationship between claimed components," that is, the term "means the claimed components must be connected in a way to perform a designated function." "In the absence of modifiers, general descriptive terms are typically construed as having their full meaning." *Id.* at 1118, 72 USPQ2d at 1006. In the patent claim at issue, "subject to any clear and unmistakable disavowal of claim scope, the term 'operatively connected' takes the full breath of its ordinary meaning, i.e., 'said tube [is] operatively connected to said cap' when the tube and cap are arranged in a manner capable of performing the function of filtering." *Id.* at 1120, 72 USPQ2d at 1008.<

Whether or not the functional limitation complies with 35 U.S.C. 112, second paragraph, is a different issue from whether the limitation is properly supported under 35 U.S.C. 112, first paragraph, or is distinguished over the prior art. A few examples are set forth below to illustrate situations where the issue of whether a functional limitation complies with 35 U.S.C. 112, second paragraph, was considered.

It was held that the limitation used to define a radical on a chemical compound as "incapable of forming a dye with said oxidizing developing agent" although functional, was perfectly acceptable because it set definite boundaries on the patent protection sought. *In re Barr*, 444 F.2d 588, 170 USPQ 33 (CCPA 1971).

In a claim that was directed to a kit of component parts capable of being assembled, the Court held that limitations such as "members adapted to be positioned" and "portions . . . being resiliently dilatible whereby said housing may be slidably positioned" serve to precisely define present structural attributes

EXHIBIT 6

RANDOM HOUSE WEBSTER'S UNABRIDGED DICTIONARY

Second Edition



RANDOM HOUSE
NEW YORK

COPYRIGHT © 1998, 1997, 1996, 1993, 1987 BY RANDOM HOUSE, INC.

All rights reserved under International and Pan-American Copyright Conventions. No part of this book may be reproduced in any form or by any means, electronic or mechanical, including photocopying, without the written permission of the publisher. All inquiries should be addressed to Random House Reference & Information Publishing, Random House, Inc., New York, NY 10022. Published in the United States by Random House, Inc., New York and simultaneously in Canada by Random House of Canada Limited.

Random House Webster's Unabridged Dictionary, Second Edition, is revised and updated edition of *The Random House Dictionary of the English Language, Second Edition, Unabridged*.

The *Random House Living Dictionary Database™*, the *Random House Living Dictionary™*, and the *Random House Living Dictionary Project™* are trademarks of Random House, Inc. *Random House* and the house design are registered trademarks of Random House, Inc.

International Phonetic Alphabet, courtesy of International Phonetic Association

Library of Congress Cataloging-in-Publication Data as of 1997:

Random House compact unabridged dictionary
Random House Webster's Unabridged dictionary.—2nd ed.
p. cm
Rev., updated ed. of: Random House compact unabridged dictionary.
Special 2nd Ed.
ISBN 0-679-45854-9.—ISBN 0-679-45853-0
I. English language—Dictionaries.
PE1625.R293 1997 423—dc21 97-17702
CIP

Trademarks

A number of entered words which we have reason to believe constitute trademarks have been designated as such. However, no attempt has been made to designate trademarks or service marks all terms or words in which proprietary rights might exist. The inclusion, exclusion, or definition of a word or term is not intended affect, or to express a judgment on, the validity or legal status of the word or term as a trademark, service mark, or other propriety term.

This book is available for special purchases in bulk by organizations and institutions, not for resale, at special discounts. Please direct your inquiries to the Random House Special Sales Department, toll free 888-591-1200 or fax 212-572-4961.

Please address inquiries about electronic licensing of reference products, for use on a network or in software or on CD-ROM, to the Subsidiary Rights Department, Random House Reference & Information Publishing, fax 212-940-7370.

New York Toronto London Sydney Auckland

This book is also sold in a special package containing the CD-ROM version 3.0 for Windows™ 95 and higher versions.

Visit the Random House Reference & Information Publishing Web site at www.randomwords.com

Typeset and printed in the United States of America.

October 1999

9 8 7 6 5 4

ISBN: 0-679-45854-9 (Hardcover/book only)

ISBN: 0-375-40383-3 (Hardcover/book and CD-ROM package)

the scales of fish, esp. of the bleak. [*< F: lit., essence of the Orient*]

es/sence of ber/gamot, bergamot (def. 2).

Es-sene (es/ēn, e sēn'), *n.* *Judaism.* a member of a Palestinian sect, characterized by asceticism, celibacy, and joint holding of property, that flourished from the 2nd century B.C. to the 2nd century A.D. — **Es-seni-an** (e sē'nē an), **Es-seni-c** (e sen'ik), *adj.*

es-sen-tial (ə sen'shəl), *adj.* 1. absolutely necessary; indispensable: *Discipline is essential in an army.* 2. pertaining to or constituting the essence of a thing. 3. not being such by its very nature or in the highest sense; natural; spontaneous: *essential happiness.* 4. *Math.* a. (of a singularity of a function of a complex variable) noting that the Laurent series at the point has an infinite number of terms with negative powers. b. (of a discontinuity) noting that the function is discontinuous and has no limit at the point. Cf. **removable** (def. 2). — **n.** 6. a basic, indispensable, or necessary element; chief point: *Concentrate on essentials rather than details.* [1300-50; ME *essencial* < ML *essenciālis* for LL *essentiālis*. See **ESSENCE**, -AL] — **es-sen-tial-ly**, *adv.* — **es-sen-tial-ness**, *n.*

— **Syn.** 1. fundamental, basic, inherent, intrinsic, vital. See **NECESSARY**. 2. **ESSENTIAL**, **INHERENT**, **INTRINSIC** refer to that which is in the natural composition of a thing. **ESSENTIAL** suggests that which is in the very essence or constitution of a thing: *Oxygen and hydrogen are essential in water.* **INHERENT** means inborn or fixed from the beginning as a permanent quality or constituent of a thing: *properties inherent in iron.* **INTRINSIC** implies belonging to the nature of a thing itself, and comprised within it, without regard to external considerations or accidentally added properties: *the intrinsic value of diamonds.* — **Ant.** 2. incidental, extraneous, extrinsic; accidental.

essen/tial am/ino ac/id, *Biochem.* any amino acid that is required by an animal for growth but that cannot be synthesized by the animal's cells and must be supplied in the diet. [1935-40]

essen/tial hyper-tension, *Pathol.* persistent high blood pressure of no known cause.

es-sen-tial-ism (ə sen'shə liz'əm), *n.* *Educ.* a doctrine that certain traditional concepts, ideals, and skills are essential to society and should be taught methodically to all students, regardless of individual ability, need, etc. Cf. **progressivism**. [1935-40; **ESSENTIAL** + -ISM] — **es-sen-tial-ist**, *n.*, *adj.*

es-sen-tial-ity (ə sen'shə al'i-tē), *n.*, *pl.* -**ities** for 2. 1. the quality of being essential; essential character. 2. an essential feature, element, or point. [1610-20; **ESSENTIAL** + -ITY]

es-sen-tial-ize (ə sen'shə līz'ē), *v.*, *-ized*, *-izing*. to extract the essence from; express the essence of. Also, *esp. Brit.*, **es-sen'tial-ise**' [1660-70; **ESSENTIAL** + -IZE]

essen/tial oil, any of a class of volatile oils obtained from plants, possessing the odor and other characteristic properties of the plant, used chiefly in the manufacture of perfumes, flavors, and pharmaceuticals. [1665-76]

es-se quam vi-dē-ri (es'sə kwām vi'dē rē; Eng. es'sē kwām vi'dār'f), *Latin.* to be rather than to seem: motto of North Carolina.

Es-se-qui-bo (es'sē kwē'bō), *n.* a river flowing from S Guyana N to the Atlantic. *l. 850 mi. (885 km) long.*

Es-sex (es'iks), *n.* 1. 2nd Earl of. See **Devereux**, **Robert**. 2. a county in SE England, 1,410,900; 1,418 sq. mi. (3,670 sq. km). 3. a town in N Maryland, near Baltimore. 39,614. 4. a town in W Vermont. 14,392.

Es-sex Jun'to (jun'tō), *U.S. Hist.* 1. a group of extreme Federalist party members from Essex county, Massachusetts. 2. any Federalist. [1795-1805, Amer.]

Es-sex table, *Carpentry.* a chart tabulating the number of board feet, to the nearest twelfth, contained in pieces of wood one inch thick and of varying standard sizes.

es-sive (es'iv), *Gram.* — **adj.** 1. noting a case, as in Finnish, whose distinctive function is to indicate a state of being. — **n.** 2. the *essive* case. [1900-05; *< Finnish essiu* < L *esse* to be + -ious -ive]

es-soin (i soin'), *n.* (in England) an excuse for nonappearance in a court of law at the prescribed time. [1300-50; ME *essoine* < AF, OF *essoigne*, *essoine*, *n.* deriv. of *essoinier* to put forward such an excuse, *v.* deriv. (with *es-*) of *soigne*, ult. *< Old Low Franconian *sunja* legal excuse, care (cf. OS *sunnea*, ON *syn* denial, Goth *sunja* truth)]

es-so-nite (es'sō nit'), *n.* *Mineral.* a variety of grossularite garnet. Also called **cinnamon stone**, **hessonite**. [1810-20; *< F < Gk hēsson* less, inferior + -ites -ite']

Es-sonne (e sōn'), *n.* a department in N France. 923,061; 699 sq. mi. (1,810 sq. km). *Cap.*: Evry.

EST, Eastern Standard Time. Also, **E.S.T.**, **e.s.t.**

-est', a suffix forming the superlative degree of adjectives and adverbs: *warmest; fastest; soonest*. [ME; OE *-est*, -ost. Cf. Gk -isto-]

-est', a native English suffix formerly used to form the second person singular indicative of verbs: *knowest; sayest; goest*. Also, **-st**. [ME; OE *-est*, -ast, -st, 2nd pers. sing. pres. indic. endings of some verbs - *e* earlier verbal ending + *-t*, by assimilation from *thū thou*'] and 2nd pers. sing. past endings of weak verbs (earlier *-es* + *-t*)]

est., 1. established. 2. estate. 3. estimate. 4. estimated. 5. estuary.

estab., established.

es-tab-lish (i stab'lish), *v.* 1. to found, institute, build, or bring into being on a firm or stable basis: *to establish a university; to establish a medical practice.* 2. to install or settle in a position, place, business, etc.: *to establish one's child in business.* 3. to show to be valid or true; prove: *to establish the facts of the matter.* 4. to

cause to be accepted or recognized: *to establish a custom; She established herself as a leading surgeon.* 5. to bring about permanently: *to establish order.* 6. to enact, appoint, or ordain for permanence, as a law; fix unalterably. 7. to make (a church) a nation or state institution. 8. *Cards.* to obtain control of (a suit) so that one can win all the subsequent tricks in it. [1325-75; ME *establisshen* < MF *establissh*, extended s. of *establi* < L *stabilire*, akin to *stabilis* STABLE'] — **es-tab-lish-a-ble**, *adj.* — **es-tab-lish-er**, *n.* — **Syn.** 1. form, organize. See **fix**. 3. verify, substantiate. 6. decree. — **Ant.** 1. abolish. 3. disprove.

estab-lished church', a church that is recognized by law, and sometimes financially supported, as the official church of a nation. Also called **state church**. Cf. **national church**. [1650-60]

es-tab-lish-ment (i stab'lish mōnt), *n.* 1. the act or an instance of establishing. 2. the state or fact of being established. 3. something established; a constituted order or system. 4. (*often cap.*) the existing power structure in society; the dominant groups in society and their customs or institutions; institutional authority (usually prec. by *the*): *The Establishment believes exploring outer space is worth any tax money spent.* 5. (*often cap.*) the dominant group in a field of endeavor, organization, etc. (usually prec. by *the*): *the literary Establishment.* 6. a household; place of residence including its furnishings, grounds, etc. 7. a place of business together with its employees, merchandise, equipment, etc. 8. a permanent civil, military, or other force or organization. 9. an institution, as a school, hospital, etc. 10. the recognition by a state of a church as the state church. 11. the church so recognized, esp. the Church of England. 12. *Archaic.* a fixed or settled income. [1476-85; 1920-25 for def. 4; ESTABLISH + -MENT]

es-tab-lish-men-tar-i-an (i stab'lish man tār'ē an), *adj.* 1. of or pertaining to an established church, esp. the Church of England, or the principle of state religion. 2. (*often cap.*) of, pertaining to, or favoring a political or social establishment. — **n.** 3. a supporter or adherent of the principle of the establishment of a church by state law; an advocate of state religion. 4. (*often cap.*) a person who belongs to or favors a political or social establishment. [1840-50; ESTABLISHMENT + -ARIAN] — **es-tab-lish-men-tar-i-an-ism**, *n.*

es-ta-fette (es'ta fet'), *n.* a mounted courier. [1785-95; *< F < It staffetta*, dim. of *staffa* stirrup < Gmc (cf. STAPES); see -ETTE]

Es-taing, d' (des tan'), Charles Hector (sharl ek-tōr'), 1729-94, French admiral.

es-ta-min (es'ta min), *n.* a worsted fabric constructed in twill weave with a rough surface. Also, **es-ta-mone** (es'ta mōn'). [1695-1705; *< F estamine* < L *stāmina*, fem. of *stāminus* made of threads. See STAMEN, -EOUS]

es-ta-min-net (es'ta mēn'et), *n.*, *pl.* -**nets** (-nē't). *French.* a bistro or small café. [1805-15]

es-tamp-je (ə stām pā'), *n.* a medieval dance and instrumental form. In several reputed sections, associated chiefly with the troubadours. [*< F, OF, deriv. of estampir* to roar, resound < Gmc; see STAMP]

es-tan-cia (e stān'shā; Sp. es tān'syā), *n.*, *pl.* -**cias** (-shēz; Sp. -syās). (in Spanish America) a landed estate or a cattle ranch. [1695-1705; *< AmerSp, Sp; dwelling*]

es-tate (i stat'), *n.*, *v.*, -**tated**, -**tating**. — **n.** 1. a piece of landed property, esp. one of large extent with an elaborate house on it: *to lease an estate in the country.* 2. *Law.* a. property or possessions. b. the legal position or status of an owner, considered with respect to property owned in land or other things. c. the degree or quantity of interest that a person has in land with respect to the nature of the right, its duration, or its relation to the rights of others. d. interest, ownership, or property in land or other things. e. the property of a deceased person, a bankrupt, etc., viewed as an aggregate. 3. *Brit.* a housing development. 4. a period or condition of life: *to attain to man's estate.* 5. a major political or social group or class, esp. one once having specific political powers, as the clergy, nobles, and commons in France or the lords spiritual, lords temporal, and commons in England. 6. condition or circumstances with reference to worldly prosperity, estimation, etc.: *social status or rank.* 7. *Obs.* pomp or state. 8. *Obs.* high social status — *v.* — **vt.** 9. *Obs.* to establish in or as in an estate. [1175-1225; ME *estat* < MF; c. Pr *estat*. See STATE] — **Syn.** 1. See **PROPERTY**.

estate/ a/gent, *Brit.* 1. the steward or manager of a landed estate. 2. a real-estate agent; realtor. [1875-80]

es-tate-bot-ting (i stāt'bot'ing, -bot'ling), *n.* a practice whereby a vineyard bottles its own wine. — **es-tate-bot'tled**, *adj.*

estate/ car', *Brit.* See **station wagon**. [1945-50]

Estates/ Gen'er-al, *French Hist.* the States-General.

estate/ tax', a tax imposed on a decedent's property, assessed on the gross estate prior to distribution to the heirs. Also called **death tax**. [1905-10]

Es-te (es'tā), *n.* a city in NE Italy; medieval fortress; ancient Roman ruins. 17,060. *Ancient, Ateste.*

es-teem (i stēm'), *v.* 1. to regard highly or favorably; regard with respect or admiration: *I esteem him for his honesty.* 2. to consider as of a certain value or of a certain type; regard: *I esteem it worthless.* 3. *Obs.* to set a value on; appraise. — **n.** 4. a favorable opinion or judgment; respect or regard: *to hold a person in esteem.* 5. *Archaic.* opinion or judgment; estimation; valuation. [1400-50; late ME *estemen*, < MF *estimer* < L *aeestimāre* to fix the value of] — **Syn.** 1. honor, revere, respect. See **APPRECIATE**. 4. favor, admiration, honor, reverence, veneration. See **RESPECT**. — **Ant.** 1. disdain.

Es-telle (i stel'), *n.* a female given name: from a Latin word meaning 'star.' Also, **Es-tel-la** (i stel'la).

Es-te-po-na (es'tē pō'nā; Eng. es'tā pō'nā), *n.* a seaport in S Spain, on the Mediterranean; resort center. 21,163.

es-ter (es'ter), *n.* *Chem.* a compound produced by the reaction between an acid and an alcohol with the elimination of a molecule of water, as ethyl acetate, C₄H₈O₂, or dimethyl sulfate, C₂H₅SO₄. [1850-55; coined by L. Gmelin (1788-1853), German chemist]

es-ter-ase (es'tā rās', -rās'), *n.* *Biochem.* any enzyme that hydrolyzes an ester into an alcohol and an acid. [1915-20; ESTER + -ASE]

es'ter gum', *Chem.* any of several hard resins produced by the esterification of a natural resin, esp. rosin, with a polyhydric alcohol, chiefly glycerol; used in the manufacture of paints, varnishes, and lacquers. [1935-40]

Es-ter-hazy (es'ter hā'zē; Fr. e ster a zē'), *n.* **Ma-rie Charles For-d'nand Walsin** (mā rē' sharl fēr dānān' vā san'), 1847-1923, French army officer who confessed forging evidence that convicted Alfred Dreyfus.

Es-ter-hazy (es'ter hā'zē; Hung. es'ter hā'zi), *n.* **Prince Mik-lós József** (mī'klōsh yō'shef), 1714-90, Hungarian patron of the arts. Also, **Es'ter-hazy**.

es-ter-ify (e ster'ə fy'), *v.*, *u.*, *-fied*, *-fying*. *Chem.* to convert into an ester. [1900-05; ESTER + -IFY] — **es-ter-ify-a-ble**, *adj.* — **es-ter-ify-ca-tion**, *n.*

Es-tes (es'tēz, -tis), *n.* a male given name.

Es'tes Park', a summer resort in N Colorado. 2703.

Esth., 1. *Bible.* Esther. 2. *Esthonia.*

Es-ther (es'ter), *n.* 1. the wife of Ahasuerus. 2. a book of the Bible bearing her name. *Abbr.:* Esth. 3. a number of prayers, visions, interpretations of dreams, etc., that are included in the Douay Bible as chapters 10-16. 4. a female given name.

es-tho-sia (es thō'shā, -zhā ə, -zē ə), *n.* capacity for sensation or feeling; sensitivity. Also, **esthōsia**. [1875-80; *< Gk aisthēsis* (see ESTHESIS) + -IA]

es-the-si-om-e-ter (es thē'si-om'ē-tor), *n.* *Med.* an instrument for measuring the degree of tactile sensibility. [ESTHESIS + -OM- + -ETER] — **es-the-si-om'o-try**, *n.*

es-the-sis (es thē'sis), *n.* sensation; feeling. Also, **esthesis**. [1850-55; *< Gk aisthēsis* sensation, perception]

es-thet-ic (es'thēt'ik), *n.* aesthetic.

es-thet-ic (es'thēt'ik), *n.* aesthetic.

es-thet-i-cal (es'thēt'ik-əl), *adj.* aesthetical. — **es-thet-i-cal-ly**, *adv.*

es-thet-i-cian (es'thēt'ish'ən), *n.* 1. aesthetician. 2. a person trained to administer facials, advise customers on makeup and the care of skin and hair, etc. Cf. **beauti-cian** (def. 1).

es-thet-i-cism (es'thēt'ə siz'əm), *n.* aestheticism.

es-thet-i-cs (es'thēt'iks), *n.* (used with a singular *v.*) aesthetics.

Es-tho-ni-a (e stō'nē ə, e stōn'yə, es thō'nē ə, -thōn'yə), *n.* Estonia.

Es-tho-ni-an (e stō'nē ən, es thō'-), *adj.*, *n.* Estonian.

Es-tienne (es tyen'), *n.* 1. Also, **Étienne**. a family of French printers, book dealers, and scholars, including esp. **Hen-ri** (ān rē'), died 1520; his son, **Ro-bert** (nō-bērt'), 1503?-59; **Hen-ri** (son of 1520), 1537-98. 2. a French printing firm founded by this family.

es-tim-a-ble (es'ti-mə-bəl), *adj.* 1. worthy of esteem; deserving respect or admiration. 2. capable of being estimated. [1425-75; late ME < MF < L *aeestimābilis*, equiv. to *aeestimāre* to ESTEEM + -ābilis -ABLE] — **es-tim-a-ble-ness**, *n.* — **es-ti-ma-bly**, *adv.* — **Syn.** 1. reputable, respectable, admirable, laudable, meritorious, excellent, good. — **Ant.** 1. contemptible.

es-ti-mate (v. es'ti-māt', n. es'ti-mit, -māt'), *v.*, -**mat-ed**, -**mat-ing**, *n.* — **vt.** 1. to form an approximate judgment or opinion regarding the worth, amount, size, weight, etc., of; calculate approximately: *to estimate the cost of a college education.* 2. to form an opinion of, judge. — **v.** 3. to make an estimate. — **n.** 4. an approximate judgment or calculation, as of the value, amount, time, size, or weight of something. 5. a judgment or opinion, as of the qualities of a person or thing. 6. a statement of the approximate charge for work to be done, submitted by a person or business firm ready to undertake the work. [1525-35; *< L aeestimātus*, ptp. of *aeestimāre* to value, estimate, see -ATE'] — **es-ti-mat'-ing-ly**, *adv.* — **es-ti-ma'tor**, *n.* — **Syn.** 1. compute, count, reckon, gauge, assess, value, evaluate, appraise. 4. valuation, calculation, appraisal.

es-ti-ma-tion (es'ti-mā'shən), *n.* 1. judgment or opinion: *In my estimation the boy is guilty.* 2. esteem; respect. 3. approximate calculation; estimate: *to make an estimation of one's expenditures.* [1325-75; ME *estimacioun* < MF < L *aeestimātiō*- (s. of *aeestimātiō*). See ESTIMATE, -ION] — **Syn.** 2. appreciation, regard, honor, veneration.

es-ti-ma-tive (es'ti-mā'tiv), *adj.* 1. capable of estimating. 2. pertaining to or based upon estimation; estimated. [1350-1400; ME < ML *aeestimātivus*. See ESTIMATE, -IVE]

es-stip-u-late (ē stip'yə lit, -lāt'), *adj.* *Bot.* exstipulate.

es-ti-val (es'ti vōl, e stī'vōl), *adj.* pertaining or appropriate to summer. [1350-1400; ME < LL *aeestivālis*, equiv. to L *aeestivus* of or relating to summer + -ālis -AL']

es-ti-vate (es'ti vāt'), *v.*, *-vated*, *-vat-ing*. 1. to spend the summer, as at a specific place or in a certain activity. 2. *Zool.* to pass the summer in a torpid condition. [1620-30; *< L aeestivātus*, ptp. of *aeestivāre* to reside

CONCISE PRONUNCIATION KEY: *act*, *cāpe*, *dāre*, *pārt*; *set*, *ēqual*; *if*, *ice*; *ox*, *ūver*, *ōrder*, *oil*, *bōok*, *bōot*, *out*, *up*, *ūrges*; *child*; *sing*; *shoe*; *thin*; *that*; *zh* as in *treasure*; *z* as in *alone*, *e* as in *system*, *i* as in *easy*; *o* as in *gallop*, *u* as in *circus*; * as in *fire* (fī'r), *hour* (hūr).
1 and *n* can serve as syllabic consonants, as in *cradle* (krād'əl), and *button* (but'n). See the full key inside the front cover.

parallel. 3. a parallel or comparison. 4. *Metaphys.* the doctrine that mental and bodily processes are concomitant, each varying with variation of the other, but that there is no causal relation of interaction between the two. [1600-10; PARALLEL + -ISM]

parallel-ist (par/ə lē'l/ist, -lə list), n. 1. a person who seeks or makes a comparison. 2. an adherent of the metaphysical doctrine of parallelism. [1785-95; PARALLEL + -IST]

parallel-ist-ic (par/ə lē'l/istik, -lə list-ik), adj. 1. of pertaining to, or of the nature of a parallelism. 2. of pertaining to the metaphysical doctrine of parallelism or to its adherents. 3. resembling, approaching, or characterized by parallelism. [1865-70; PARALLELIST + -IC]

parallel-ize (par/ə lē'l/iz, -lə lē'l/iz), v.t., -ized, -izing. 1. to make parallel; place so as to be parallel. 2. to draw a parallelism or analogy between. Also, esp. *Brit.*, **par/ə lē'l/iz**. [1600-10; < Gk *parallēlizein*. See PARALLEL, -IZE] —**par/ə lē'l/iz-a-tion**, n.

par/allel mo-tion, a mechanism arranged so as to impart rectilinear motion to a rod connected to a lever that moves through an arc. [1820-30]

par/allel of al-titude, *Astron.* almucaantar. [1695-1705]

par/allel of lat-itude, parallel (def. 9). [1660-70]

par/allel-o-gram (par/ə lē'l/ə gram), n. a quadrilateral having both pairs of opposite sides parallel to each other. [1560-70; < LL *parallēlogrammum* < Gk *parallēlogrammōn*. See PARALLEL, -O-, -GRAM] —**par/allel-o-gram-mat-ic**, (par/ə lē'l/ə grə mat'ik), **par/allel-o-gram-mat-i-cal**, adj.



parallel-ogram law, *Math. Physics.* a rule for adding two vectors, as forces (**parallel-ogram of force**), by placing the point of application of one at the point of origin of the other and obtaining their sum by constructing the line connecting the two remaining end points, the sum being the diagonal of the parallelogram whose adjacent sides are the two vectors.

par/allel pos-tulate, *Geom.* the axiom in Euclidean geometry that only one line can be drawn through a given point so that the line is parallel to a given line that does not contain the point. Also called **parallel axiom**.

par/allel projec-tion, *Geom.* a projection from one plane to a second plane in which the lines joining points on the first plane and corresponding images are parallel.

par/allel rul-ers, a pair of straightedges connected by two pivoted crosspieces of equal length so as to be parallel at all times, used for various navigational purposes, esp. for transferring the bearing of a plotted course to a compass rose. [1695-1705]

par/allel sail-ing, sailing along a parallel of latitude. [1700-10]

par/allel top, See under **parallel** (def. 16).

par/allel-veined (par/ə lē'l/vēnd, -ləl-), adj. *Bot.* having closely spaced longitudinal veins, as the leaves of most monocotyledonous plants. Cf. **net-veined**. [1860-65; PARALLEL + VEIN + -ED]

par/alog-ism (pə ral/ə jiz/əm), n. *Logic.* 1. argument violating principles of valid reasoning. 2. a conclusion reached through such argument. [1565-65; < LL *paralogismus* < Gk *paralogismōs*. See PARA-, LOGO-, -ISM] —**par/alog-ist**, n. —**par/alog-ist-ic**, adj.

par/alog-ize (pə ral/ə jiz/), v.t., -gized, -gizing. to draw conclusions that do not follow logically from a given set of assumptions. Also, esp. *Brit.*, **par/alog-ize**. [1590-1600; < ML *paralogizāre* < Gk *paralogizesthai* to reason falsely, equiv. to *paralogōs* (see PARA-, LOGOS) + *-izesthai* -IZE]

par/aly-sis (pə ral/ə sis), n., pl. -ses (-sēz/). 1. *Pathol.* a. a loss or impairment of voluntary movement in a body part, caused by injury or disease of the nerves, brain, or spinal cord. b. a disease characterized by this, esp. palsy. 2. a state of helpless stoppage, inactivity, or inability to act. The strike caused a paralysis of all shipping. [bef. 1150; < L < Gk *paralysis*, equiv. to *paraly-* var s. of *paralyein* to loosen (i.e., disable) on one side (para- PARA- + *lyein* to loosen) + *-sis* -SIS; r. ME *paralis(e)* < OF < L, as above; r. late OE *paralisan* (acc.) < L, as above; cf. **PALSY**]

par/alyt-ic (par/ə lit'ik), n. 1. a person affected with paralysis. —adj. 2. affected with or subject to paralysis. 3. pertaining to or of the nature of paralysis. [1300-50; ME *paralitic* < L *paralyticus* < Gk *paralytikōs*, equiv. to *paraly-* (see PARALYSIS) + *-tikos* -TIC] —**par/alyt-i-cal-ly**, adv.

par/alyz-e (par/ə liz/), v.t., -lyzed, -lyzing. 1. to affect with paralysis. 2. to bring to a condition of helpless stoppage, inactivity, or inability to act. The strike paralyzed communications. Also, esp. *Brit.*, **par/alyz-e**. [1795-1805; back formation from PARALYSIS, modeled on *analyze*] —**par/alyz-a-tion**, n. —**par/alyz-er**, n. —**par/alyz-ing-ly**, adv. —Syn. 2. See **shock**.

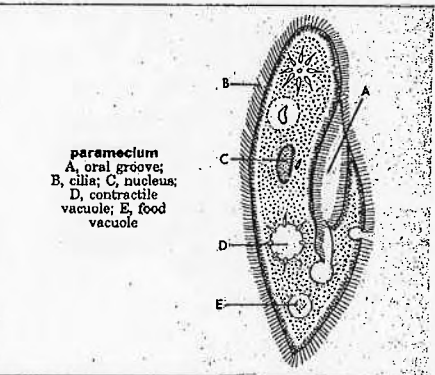
par/a-mag-net (par/ə mag'nit, par/ə mag'/), n. *Physics.* a body or substance that, placed in a magnetic field, possesses magnetization in direct proportion to the field strength; a substance in which the magnetic moments of the atoms are not aligned. Cf. **antiferromagnetic**, **diamagnetic**, **ferromagnetic**. [1905-10; back formation

from *paramagnetic*; see PARA-, MAGNET] —**par/a-mag-net-ism**, n. —**par/a-mag-net-ic** (par/ə mag net'ik), adj.

Par/a-mar-i-bo (par/ə mar'ə bō/), n. a seaport in and the capital of Suriname, in NE South America. 150,000.

Par/am-at-man (par/ə māt'man), n. *Hinduism.* absolute Atman. [*< Skt paramātman supreme self*]

par/a-mat-ta (par/ə mat'ə), n. a light, twilled dress fabric, having a silk or cotton warp and a woolen weft. Also, **Parramatta**. [1825-35; named after Parramatta, town in New South Wales]



paramecium
A, oral groove;
B, cilia; C, nucleus;
D, contractile
vacuole; E, food
vacuole

par/a-me-ci-um (par/ə mē'shē əm, -shəm, -sē əm), n., pl. -cia (-shē ə, -sē ə). any ciliated freshwater protozoan of the genus *Paramecium*, having an oval body and a long, deep oral groove. [1745-55; < NL < Gk *paramēkēs* oblong, oval + NL *-ium* n. suffix; see -IUM]

par/a-med-ic (par/ə med'ik), n. a person who is trained to assist a physician or to give first aid or other health care in the absence of a physician, often as part of a police, rescue, or firefighting squad. [1950-55, *Amer.*; PARA- + MEDIC]

par/a-med-ic (n. par/ə med'ik, par/ə med'/; adj. par/ə med'ik), n. 1. *Mil.* a medic in the paratroops. 2. a doctor who parachutes into remote areas to give medical care. —adj. 3. of or pertaining to a paramedic or to paramedics. [1950-55, *Amer.*; PARA- + MEDIC]

par/a-med-i-cal (par/ə med'i kal), adj. related to the medical profession in a secondary or supplementary capacity. [1920-25; PARA- + MEDICAL]

para-ment (par/ə ment), n., pl. **para-ments**, **para-ment** (par/ə men'tə), n. 1. a decoration for a room, as a tapestry. 2. an ecclesiastical vestment. [1350-1400; ME < LL *paramentum* an ornament, equiv. to *parā(re)* to adorn (L to PREPARE) + *-mentum* -MENT]

para-met-er (pə ram'i tər), n. 1. *Math.* a constant or variable term in a function that determines the specific form of the function but not its general nature, as in $f(x) = ax$, where a determines only the slope of the line described by $f(x)$. b. one of the independent variables in a set of parametric equations. 2. *Statistics.* a variable entering into the mathematical form of any distribution such that the possible values of the variable correspond to different distributions. 3. *Computers.* a variable that must be given a specific value during the execution of a program or of a procedure within a program. 4. Usually, **parameters**, limits or boundaries; guidelines: the basic parameters of our foreign policy. 5. characteristic or factor; aspect; element: a useful parameter for judging long-term success. [1650-60; < NL *parametrum*. See PARA-, -METER] —**par/a-met-ric** (par/ə mē'trik), **par/a-met-ri-cal**, adj. —Usage. 4. Some object strongly to the use of PARAMETER in these newer senses. Nevertheless, the criticized uses are now well established both in educated speech and in edited writing.

par/a-met-er-ize (pə ram'i tər/iz), v.t., -ized, -izing. to describe (a phenomenon, problem, curve, surface, etc.) by the use of parameters. Also, **parameterize**; esp. *Brit.*, **par/a-met-er-ize**. [1935-40; PARAMETER + -IZE] —**par/a-met-er-iz-a-tion**, n.

par/a-met-er-iz-a-tion, n.

par/a-met-ric am-plifier, *Electronics.* a device, as an electron tube or transistor, that amplifies a high-frequency input signal by sinusoidally varying the reactance of the circuit. [1955-60]

par/a-met-ric equa-tion, *Math.* one of two or more equations expressing the location of a point on a curve or surface by determining each coordinate separately. [1905-10]

par/a-met-ri-ze (pə ram'i triz/), v.t., -trized, -trizing. parameterize. Also, esp. *Brit.*, **par/a-met-ri-ze**.

par/a-mil-i-tary (par/ə mil'i tər'ē), adj., n., pl. -taries. —adj. 1. noting or pertaining to an organization operating as, in place of, or as a supplement to a regular military force: a paramilitary police unit. —n. 2. Also, **par/a-mil-i-tar-ist** (par/ə mil'i tər'ist), a person employed in such a force. [1930-35; PARA- + MILITARY]

par/a-mit-ta (pə rum'i tā), n. *Buddhism.* any of the practices prescribed for one aspiring to nirvana. [*< Skt and Pali pāramitā perfection*]

par/a-mne-sia (par'am nē'zha), n. 1. *Psychiatry.* a distortion of memory in which fact and fantasy are confused. 2. the inability to recall the correct meaning of a word. [1885-90; < NL; see PARA-, AMNESIA]

par/a-mo (par/ə mō, pār'ə-), n., pl. -mos. a high, cold plateau of South America. [1750-60; < AmerSp; Sp *páramo* barren plain; presumably of pre-L orig.]

par/a-morph (par/ə mōrf'), n. *Mineral.* a pseudomorph formed by a change in crystal structure but not in

chemical composition. Also called **allomorph**. [1875-80; PARA- + -MORPH] —**par/a-mor-phic**, **par/a-mor-phous**, adj.

par/a-mor-phine (par/ə mōr'fēn), n. *Chem.* theba-ine. [PARA- + MORPHINE]

par/a-mor-phism (par/ə mōr'fiz əm), n. 1. the process by which a paramorph is formed. 2. the state of being a paramorph. [1865-70; PARA- + -MORPHISM]

par/a-mount (par/ə maunt/), adj. 1. chief in importance or impact; supreme; preeminent: a point of paramount significance. 2. above others in rank or superiority; superior in power or jurisdiction. —n. 3. a supreme ruler; overlord. [1525-35; < AF *paramont* above, equiv. to *par* FER- + *a mont* < L *ad montem* to the mountain, hence, in OF, upward, above; see AD-, MOUNT] —**par/a-mount-ly**, adv. —Syn. 1. See **dominant**. —Ant. 1. unimportant.

Par/a-mount (par/ə maunt/), n. a city in SW California, near Los Angeles. 96,407.

par/a-mour (par/ə mōōr/), n. 1. an illicit lover, esp. of a married person. 2. any lover. [1250-1300; ME, from the phrase *par amour* by or through love < OF]

Par/a-mus (pə ram'əs), n. a city in NE New Jersey. 26,474.

par/a-myx-o-vi-rus (par/ə mik'sə vi'ras, -mik'sə-vi-rə/), n., pl. -ruses. any of various RNA-containing viruses that are similar to but larger than the myxoviruses, including the viruses that cause mumps, measles, parainfluenza, and Newcastle disease. [1960-65; PARA- + MYXOVIRUS]

Par/a-ná (par/ə nā/), *Port.* pá'rã ná/), n. 1. a river in central South America, flowing from S Brazil along the SE boundary of Paraguay and through E Argentina into the Rio de la Plata. 2650 mi. (3300 km) long. 2. a city in E Argentina, on the Paraná River; the capital of Argentina 1852-61. 169,581.

Par/a-na-guá (pə nã'güá/), n. a seaport in S Brazil. 65,178.

par/a-na-sal (par/ə nã'sæl), adj. *Anat.* situated near the nasal cavities. [1905-10; PARA- + NASAL]

par/ang (pär'äng), n. a large, heavy knife used as a tool or a weapon in Malaysia and Indonesia. [1850-55; < Malay]

par/a-ni-tro-ph-e-nol (par/ə nit'rə fē'nól, -nol), n. See under **nitrophenol** (def. 2).

par/a-noi-a (par/ə noi'ə), n. 1. *Psychiatry.* a mental disorder characterized by systematized delusions and the projection of personal conflicts, which are ascribed to the supposed hostility of others, sometimes progressing to disturbances of conscience and aggressive acts believed to be performed in self-defense or as a mission. 2. baseless or excessive suspicion of the motives of others. Also, **par/a-noe-a** (par/ə nō'ə). [1805-15; < NL < Gk *paranoia* madness. See PARA-, NOUS, -IA]

par/a-noid (par/ə noi'd), adj. 1. of, like, or suffering from paranoia. —n. 2. a person suffering from paranoia. Also, **par/a-noi-ac** (par/ə noi'ak, -ik), **par/a-noe-ac** (par/ə nō'ak, -ik). [1900-05; PARANOIA + -OID, with base and suffix merged, perh. by haplogly from the expected *paranoioid]

par/a-nor-mal (par/ə nōr'məl), adj. of or pertaining to the claimed occurrence of an event or perception without scientific explanation, as psychokinesis, extrasensory perception, or other purportedly supernatural phenomena. [1915-20; PARA- + NORMAL] —**par/a-nor-mal-ly**, adv.

par/a-nthro-pus (pə ran'thrə pəs, par'ən thrō'/), n., pl. -puses for 1. (sometimes cap.) a member of the former genus *Paranthropus*. 2. (cap., italics) a former genus of fossil hominids whose members have now been assigned to the proposed species *Australopithecus robustus*. [*< NL (1938) < Gk par-* PAR- (in the sense "near") + *ánthrōpos* man]

par/a-ny-mph (par/ə nimf'), n. 1. a groomsmen or a bridesmaid. 2. (in ancient Greece) a. a friend who accompanied the bridegroom when he came to bring home the bride. b. the bridesmaid who escorted the bride to the bridegroom. [1585-95; < LL *paranympus* < Gk *paránympnos* (masc. and fem.) groomsmen, bridesmaid, lit., person beside the bride. See PARA-, NYMPH]

par/a-pa-re-sis (par/ə pə rē'sis, -par/ə sis), n. *Pathol.* partial paralysis, esp. of the lower limbs. [*< NL*; see PARA-, PARESIS]

par/a-pet (par/ə pit, -pet/), n. 1. *Fort.* a. a defensive wall or elevation, as of earth or stone, in a fortification. See **diag.** under **bastion**. b. an elevation raised above the main wall or rampart of a permanent fortification. 2. any low protective wall or barrier at the edge of a balcony, roof, bridge, or the like. [1675-85; < *It* *parapetto*, equiv. to *para-* PARA- + *petto* chest, breast < *L pectus*] —**par/a-pet-ed**, adj. —**par/a-pet-less**, adj.

par/aph (par'ef, pə raf/), n. a flourish made after a signature, as in a document, originally as a precaution against forgery. [1350-1400; ME *paraf* < *It* *parafo* or MF *paraphe* paragraph mark (by syncope; see PARAGRAPH)]

par/a-phase am-plifier (par/ə fāz/), *Electronics.* an amplifier that produces a push-pull output from a single input. [PARA- + PHASE]

par/a-phe-net-i-dine (par/ə fə net'i dēn', -dīn), n. *Chem.* See under **phenetidine**.

par/a-pher-na-lia (par/ə fər nāl'yə, -fə nāl'yə/), n. 1. (sometimes used with a singular *u*) equipment, apparatus, or furnishing used in or necessary for a particular activity: a skier's paraphernalia. 2. (used with a plural *u*) personal belongings. 3. (used with a singular *u*) *Law.* the personal articles, apart from dower, reserved by law to a married woman. [1470-80; < ML *paraphernalia* (bona) a bride's goods, beyond her dowry, equiv. to LL *parapherna* (a) a bride's property (< Gk *parápherna*, equiv. to *para-* PARA- + *pherna* dowry, deriv. of *pherein* to BEAR + *-a* neut. pl. n. suffix) + *līa*, n.

CONCISE ETYMOLOGY KEY: <, descended or borrowed from; >, whence; b., blend of; blend; c., cognate with; cf., compare; deriv., derivative; equiv., equivalent; imit., imitative; obl., oblique; r., replacing; s., stem; sp., spelling, spelled; resp., respelling, respelled; trans., translation; ?, origin unknown; †, unattested; ‡, probably earlier than. See the full key inside the front cover.

EXHIBIT 7

Microsoft Press
**Computer
Dictionary**

Third Edition

Microsoft Press

PUBLISHED BY
Microsoft Press
A Division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Copyright © 1997 by Microsoft Corporation

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Cataloging-in-Publication Data
Microsoft Press Computer Dictionary. -- 3rd ed.

p. cm.

ISBN 1-57231-446-X

1. Computers--Dictionaries. 2. Microcomputers--Dictionaries.

I. Microsoft Press.

QA76.15.M54 1997

004'.03--dc21

97-15489

CIP

Printed and bound in the United States of America.

2 3 4 5 6 7 8 9 QMQM 2 1 0 9 8 7

Distributed to the book trade in Canada by Macmillan of Canada, a division of Canada Publishing Corporation.

A CIP catalogue record for this book is available from the British Library.

Microsoft Press books are available through booksellers and distributors worldwide. For further information about international editions, contact your local Microsoft Corporation office. Or contact Microsoft Press International directly at fax (425) 936-7329.

Macintosh, Power Macintosh, QuickTime, and TrueType are registered trademarks of Apple Computer, Inc. Intel is a registered trademark of Intel Corporation. DirectInput, DirectX, Microsoft, Microsoft Press, MS-DOS, Visual Basic, Visual C++, Win32, Win32s, Windows, Windows NT, and XENIX are registered trademarks and ActiveMovie, ActiveX, and Visual J++ are trademarks of Microsoft Corporation. Java is a trademark of Sun Microsystems, Inc. Other product and company names mentioned herein may be the trademarks of their respective owners.

Acquisitions Editor: Kim Fryer

Project Editor: Maureen Williams Zimmerman, Anne Taussig

Technical Editors: Dail Magee Jr., Gary Nelson, Jean Ross, Jim Fuchs, John Conrow, Kurt Meyer, Robert Lyon, Roslyn Lutsch

nection can. The parallel interface is preferred in the IBM PC world because its cabling is more standardized than that of the serial interface and because the MS-DOS operating system assumes that the system printer is attached to the parallel port. *See also* parallel interface. *Compare* serial printer.

parallel processing \pâr-ə-ləl pros-ēs-ēng\ *n.* A method of processing that can run only on a computer that contains two or more processors running simultaneously. Parallel processing differs from multiprocessing in the way a task is distributed over the available processors. In multiprocessing, a process might be divided up into sequential blocks, with one processor managing access to a database, another analyzing the data, and a third handling graphical output to the screen. Programmers working with systems that perform parallel processing must find ways to divide a task so that it is more or less evenly distributed among the processors available. *Compare* coprocessor, multiprocessing.

parallel server \pâr-ə-ləl sər-vər\ *n.* A computer system that implements some form of parallel processing to improve its performance as a server. *See also* SMP server.

parallel transmission \pâr-ə-ləl tranz-mish-ən\ *n.* The simultaneous transmission of a group of bits over separate wires. With microcomputers, parallel transmission refers to the transmission of 1 byte (8 bits). The standard connection for parallel transmission is known as the Centronics interface. *See also* Centronics parallel interface. *Compare* serial transmission.

parameter \pər-am-ə-tər\ *n.* In programming, a value that is given to a variable, either at the beginning of an operation or before an expression is evaluated by a program. Until the operation is completed, a parameter is effectively treated as a constant value by the program. A parameter can be text, a number, or an argument name assigned to a value that is passed from one routine to another. Parameters are used as a means of customizing program operation. *See also* argument, pass by address, pass by value, routine.

parameter-driven \pər-am-ə-tər-driv-ən\ *adj.* Of, pertaining to, or being a program or an operation whose character or outcome is determined by the values of the parameters that are assigned to it.

parameter passing \pər-am-ə-tər pas-ēng\ *n.* In programming, the substitution of an actual parameter value for a formal parameter when a procedure or function call is processed.

parameter RAM \pər-am-ə-tər ram, R-A-M\ *n.* A few bytes of battery-backed CMOS RAM on the motherboards of Apple Macintosh computers. Information about the configuration of the system is stored in parameter RAM. *Acronym:* PRAM (P'ram, P'R-A-M', pram). *See also* CMOS RAM. *Compare* CMOS (definition 2).

PARC \pärk, P'A-R-C\ *n.* *See* Xerox PARC.

parent/child \pâr-ənt-chīld\ *adj.* **1.** Pertaining to or constituting a relationship between processes in a multitasking environment in which the parent process calls the child process and most often suspends its own operation until the child process aborts or is completed. **2.** Pertaining to or constituting a relationship between nodes in a tree data structure in which the parent is one step closer to the root (that is, one level higher) than the child.

parity \pâr-ə-tē\ *n.* The quality of sameness or equivalence, in the case of computers usually referring to an error-checking procedure in which the number of 1s must always be the same—either even or odd—for each group of bits transmitted without error. If parity is checked on a per-character basis, the method is called vertical redundancy checking, or VRC; if checked on a block-by-block basis, the method is called longitudinal redundancy checking, or LRC. In typical modem-to-modem communications, parity is one of the parameters that must be agreed upon by sending and receiving parties before transmission can take place. Types of parity are shown in the following table. *See also* parity bit, parity check, parity error.

Type	Description
Even parity	The number of 1s in each successfully transmitted set of bits must be an even number.
Odd parity	The number of 1s in each successfully transmitted set of bits must be an odd number.
No parity	No parity bit is used.
Space parity	A parity bit is used and is always set to 0.
Mark parity	A parity bit is used and is always set to 1.

priorities that indicate how soon they must be transmitted. *See also* interrupt.

Priority Frame \prī-ōr'ə-tē frām\ *n.* A telecommunications protocol developed by Infonet and Northern Telecom, Inc., designed to carry data, facsimile, and voice information.

privacy \prī'və-sē\ *n.* The concept that a user's data, such as stored files and e-mail, is not to be examined by anyone else without that user's permission. A right to privacy is not generally recognized on the Internet. Federal law protects only e-mail in transit or in temporary storage, and only against access by Federal agencies. Employers often claim a right to inspect any data on their systems. To obtain privacy, the user must take active measures such as encryption. *See also* encryption, PGP, Privacy Enhanced Mail. *Compare* security.

Privacy Enhanced Mail \prī'və-sē en-hansd' māl\ *n.* An Internet standard for e-mail systems that use encryption techniques to ensure the privacy and security of messages. *Acronym:* PEM (P'E-M'). *See also* encryption, standard. *Compare* PGP.

Private Branch Exchange \prī'vət branch' eks-chānj\ *n.* *See* PBX.

private channel \prī'vət chan'əl\ *n.* In Internet relay chat (IRC), a channel reserved for the use of a certain group of people. Private channel names are hidden from view by the public at large. *Also called* secret channel. *See also* IRC.

Private Communications Technology \prī'vət kə-myōō'nə-kā'shənz tek-nol'ə-jē\ *n.* A specification designed to secure general-purpose business and personal communications on the Internet, and including features such as privacy, authentication, and mutual identification.

private folders \prī'vət fōl'dərz\ *n.* In a shared network environment, those folders on a user's computer that are not currently accessible by other users on the network. *Compare* public folders.

private key \prī'vət kē\ *n.* One of two keys in public key encryption. The user keeps the private key secret and uses it to encrypt digital signatures and to decrypt received messages. *See also* public key encryption. *Compare* public key.

private line \prī'vət līn\ *n.* *See* dedicated line (definition 1).

privatization \prī'və-tə-zā'shən\ *n.* Generally, the process of turning something over from gov-

ernment to commercial industry control. In the context of computer science and the Internet, the term refers to the government's turning over of various Internet backbones to private industry. For example, control of NSFnet was passed from the government to private business in 1992.

privileged instruction \priv'ə-ləjd in-struk'shən\ *n.* An instruction (usually a machine instruction) that can be executed only by the operating system. Privileged instructions exist because the operating system needs to perform certain operations that applications should not be allowed to perform; therefore, only the operating-system routines have the necessary privilege to execute these particular instructions.

privileged mode \priv'ə-ləjd mōd\ *n.* A mode of execution, supported by the protected mode of the Intel 80286 and higher microprocessors, in which software can carry out restricted operations that manipulate critical components of the system, such as memory and input/output ports (channels). Application programs cannot be executed in privileged mode; the heart (kernel) of the OS/2 operating system can be, as can the programs (device drivers) that control devices attached to the system.

privileges \priv'ə-lə-jəz, priv'lə-jəz\ *n.* *See* access privileges.

PRN \P'R-N\ *n.* The logical device name for *printer*. A name reserved by the MS-DOS operating system for the standard print device. PRN usually refers to a system's first parallel port, also known as LPT1.

probability \prob'ə-bil'ə-tē\ *n.* The likelihood that an event will happen, which can often be estimated mathematically. In mathematics, statistics and probability theory are related fields. In computing, probability is used to determine the likelihood of failure or error in a system or device.

problem solving \pro'bləm sol'veng\ *n.* 1. The process of devising and implementing a strategy for finding a solution or for transforming a less desirable condition into a more desirable one. 2. An aspect of artificial intelligence wherein the task of problem solving is performed solely by a program. *See also* artificial intelligence.

procedural language \prə-sē'jər-əl lang'wəj\ *n.* A programming language in which the basic pro-

EXHIBIT 8

DICTIONARY

OF

**COMPUTER
SCIENCE,
ENGINEERING,
AND TECHNOLOGY**

EDITOR-IN-CHIEF

Phillip A. Laplante



CRC Press

Boca Raton London New York Washington, D.C.

Library of Congress Cataloging-in-Publication Data

Dictionary of computer science, engineering, and technology / edited by Phillip Laplante.

p. cm.

ISBN 0-8493-2691-5 (alk. paper)

1. Computer science—Dictionaries. 2. Engineering—Dictionaries. 3. Technology—Dictionaries. I. Laplante, Phillip A.

QA76.15.D5258 2000

004'.03—dc21

00-052882

This book contains information obtained from authentic and highly regarded sources. Reprinted material is quoted with permission, and sources are indicated. A wide variety of references are listed. Reasonable efforts have been made to publish reliable data and information, but the author and the publisher cannot assume responsibility for the validity of all materials or for the consequences of their use.

Neither this book nor any part may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, microfilming, and recording, or by any information storage or retrieval system, without prior permission in writing from the publisher.

All rights reserved. Authorization to photocopy items for internal or personal use, or the personal or internal use of specific clients, may be granted by CRC Press LLC, provided that \$.50 per page photocopied is paid directly to Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923 USA. The fee code for users of the Transactional Reporting Service is ISBN0-8493-2691-5/01/\$0.00+\$.50. The fee is subject to change without notice. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

The consent of CRC Press LLC does not extend to copying for general distribution, for promotion, for creating new works, or for resale. Specific permission must be obtained in writing from CRC Press LLC for such copying.

Direct all inquiries to CRC Press LLC, 2000 N.W. Corporate Blvd., Boca Raton, Florida 33431.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation, without intent to infringe.

© 2001 by CRC Press LLC

No claim to original U.S. Government works

International Standard Book Number 0-8493-2691-5

Library of Congress Card Number 00-052882

Printed in the United States of America 1 2 3 4 5 6 7 8 9 0

Printed on acid-free paper

ME

which acts to enhance changes (temporal or spatial) in the stimulus.

lattice (1) a directed acyclic graph which has one designated root node, the top node, and a designated bottom node. The bottom node has the property that no edges leave this node. Furthermore, any traversal that starts at the root node will always end at this bottom node. *See also* semi-lattice.

(2) a point lattice generated by taking integer linear combinations of a set of basis vectors.

lattice vector quantization a structured vector quantizer where the reproduction vectors are chosen from a highly regular geometrical structure known as a "lattice". The method is employed mainly because of the reduction in storage capacity obtained (compared to optimal vector quantization).

lattice VQ *See* lattice vector quantization.

L-attributed grammar an attribute grammar whose attributes may be computed by a left to right traversal of the source program. An attribute grammar must be L-attributed for the attributes to be computable during a parse that processes the input from left to right (as most parsers do). Synthesized attributes are always L-attributed.

layer assignment given a set of trees in the plane, each interconnecting the terminals of a net, an assignment of a routing layer to each segment of each tree so that the resulting wiring layout is legal under the routing model.

layered queuing model an extension of queuing models that allows reasoning about client/server architectures and the performance impacts of resource requests at different layers.

lazy evaluation an optimization technique applied to the execution of an algorithm by which the actual computation specified by the algorithm is deferred until the result is required. This can mean that many computations need never be performed at all. *See* eager evaluation.

LCD *See* liquid-crystal display.

LCFS *See* last-come-first-serve.

leading-zeros anticipator a hardware unit that predicts the position of the first non-0 digit in an expected sequence of digits and encodes this position as a number; the prediction may be slightly off and so may require a small adjustment. Most commonly used to predict the length of a normalization shift in a floating-point addition and hence to speed up the process by carrying the shifting concurrently with the significant addition. *See also* leading-zeros detector.

leading-zeros detector a hardware unit that detects the position of the first non-0 digit in a sequence of digits and encodes this position as a number. Most commonly used to determine the distance for a normalization shift following a floating-point addition. *See also* leading-zeros anticipator.

leading-zeros predictor *See* leading-zeros anticipator.

leaf node a node in a tree which has the property that there are no arcs out of it to other nodes. *Contrast with* tree node and internal node. *See also* root node. In a tree with only one node (a special case), the single node is both the root node and a leaf node.

leaf procedure a procedure that does not call another procedure.

learnability the capability of the software product to enable the user to learn its application.

learning (1) generally, any scheme whereby experience or past actions and reactions are automatically used to change parameters in an algorithm.

(2) in neural networks, the collection of learning rules or laws associated with each processing element. Each learning law is responsible for adapting the input-output behavior of the processing element transfer function over a period of time in response to the input signals that influence the processing element. This adaptation is usually obtained by modification of the values of variables (weights) stored in the processing element's local memory.

M

M (mega) abbreviation for 1,048,576 (not for 1 million).

MAC See medium access control, mandatory access control.

MAC address synonym for an IEEE 802 address.

Mach band a perceived overshoot on the light side of an edge and an undershoot on the dark side of the edge. The *Mach band* is an artifact of the human visual system and not actually present in the edge. See also brightness, simultaneous contrast.

machine code (1) the native representation of a program for a specific machine architecture.

(2) source code in assembly language.

(3) an internal representation of the target machine instructions in a compiler. Often called machine language or object code.

machine epsilon the relative error when a number is rounded to the closest machine-representable number.

machine independent pertaining to software that can be executed on many platforms. Compare with portability.

machine interference the idle time experienced by any one machine in a multiple-machine system that is being serviced by an operator (or robot) and is typically measured as a percentage of the total idle time of all the machines in the systems to the operator (or robot) cycle time.

machine language the set of legal instructions to a machine's processor, expressed in binary notation. See machine code.

machine learning (1) In knowledge discovery, machine learning is most commonly used to

mean the application of induction algorithms, which is one step in the knowledge discovery process. Machine learning is the field of scientific study that concentrates on induction algorithms and on other algorithms that can be said to "learn".

(2) the component of artificial intelligence that deals with the algorithms that improve with experience.

machine simulation that aspect of code generation that determines the preconditions for and postconditions of executing a particular instruction or instruction sequence.

machine translation translating a text in one natural language to another natural language by computer.

machine vision See robot vision.

macro a construct that specifies a source-to-source translation. A macro definition specifies the translation. When an instance of the macro occurs later in the program, it is expanded according to the definition. A macro definition may specify zero or more macro parameters that are replaced by text specified at the place the macro is used. Often the syntax and semantics of macros are substantially different from the syntax and semantics of the language in which they are used. The power of a system of macros may be as simple as straight textual substitution or as complex as a form of symbolic evaluation. See also macroprogram.

macro cycle the main repetitive set of activities that are performed in the main cycle of evolutionary life cycles, such as spiral.

macro development all the strategic activities related to the design of system architecture and system-level test.

macroinstruction (1) the lowest level of user-programmable computer instruction. See opcode.

(2) a shorthand for a number of language instructions or in integrated environments. In the latter, the definition of macros is a way to make shorter the execution and the writing of repeti-

EXHIBIT 9

LEARNING FROM DATA

Concepts, Theory, and Methods

Vladimir Cherkassky
Filip Mulier



A WILEY-INTERSCIENCE PUBLICATION

JOHN WILEY & SONS, INC.

New York / Chichester / Weinheim / Brisbane / Singapore / Toronto

This book is printed on acid-free paper.

Copyright © 1998 by John Wiley & Sons, Inc. All rights reserved.

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (508) 750-8400, fax (508) 750-4744. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 605 Third Avenue, New York, NY 10158-0012, (212) 850-6011, fax (212) 850-6008, E-mail: PERMREQ@WILEY.COM.

Library of Congress Cataloging-in-Publication Data:

Cherkassky, Vladimir S.

Learning from data: concepts, theory, and methods / Vladimir Cherkassky and Filip Mulier.

p. cm.

“A Wiley-Interscience publication.”

Includes bibliographical references and index.

ISBN 0-471-15493-8 (cloth: alk. paper)

1. Adaptive signal processing. 2. Machine learning. 3. Neural networks (Computer science) 4. Fuzzy systems. I. Mulier, Filip. II. Title.

TK5102.9.C475 1998

006.3'1--dc21

97-43019

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

System The system produces an output value y for every input vector \mathbf{x} according to the fixed conditional density $p(y | \mathbf{x})$, which is also unknown. Note that this description includes the specific case of a deterministic system where $y = f(\mathbf{x})$ as well as the regression formulation of $y = f(\mathbf{x}) + \epsilon$ where ϵ is random noise with zero mean. Real systems rarely have truly random outputs; however, they often have unmeasured inputs (Fig. 1.1). Statistically the effect of these changing unobserved inputs on the output of the system can be characterized as random and represented as a probability distribution.

Learning Machine In the most general case, the learning machine is capable of implementing a set of functions $f(\mathbf{x}, \omega)$, $\omega \in \Omega$, where Ω is a set of abstract parameters used only to index the set of functions. In this formulation the set of functions implemented by the learning machine can be any set of functions, chosen a priori, before the formal inference (learning) process is begun. Let us look at some simple examples of learning machines and how they fit this formal description. The examples chosen are all solutions to the regression problem, which is only one of the four most common learning tasks (Section 2.1.2). The examples illustrate the notion of a set of functions (of a learning machine) and not the mechanism by which the learning machine chooses the best approximating function from this set.

Example 2.1 Parametric Regression (Fixed Degree Polynomial)

In this example the set of functions is specified as a polynomial of fixed degree, and the training data have a single predictor variable ($x \in \mathfrak{R}^1$). The set of functions implemented by the learning machine is

$$f(x, \mathbf{w}) = \sum_{i=0}^{M-1} w_i x^i \quad (2.1)$$

where the set of parameters Ω take the form of vectors $\mathbf{w} = [w_0, \dots, w_{M-1}]$ of fixed length M .

Example 2.2 Semiparametric Regression (Polynomial of Arbitrary Degree)

One way to provide a wider class of functions for the learning machine is to remove the restriction of fixed polynomial degree. The degree of the polynomial now becomes another parameter that indexes the set of functions

$$f_m(x, \mathbf{w}_m) = \sum_{i=0}^{m-1} w_i x^i \quad (2.2)$$

Here the set of parameters Ω take the form of vectors $\mathbf{w}_m = [w_0, \dots, w_{m-1}]$, which have an arbitrary length m .

$$(\mathbf{x}_i, y_i), \quad (i = 1, \dots, n) \quad (2.6)$$

The quality of an approximation produced by the learning machine is measured by the loss $L(y, f(\mathbf{x}, \omega))$ or discrepancy between the output produced by the system and the learning machine for a given point \mathbf{x} . By convention, the loss takes on nonnegative values, so that large positive values correspond to poor approximation. The expected value of the loss is called the *risk functional*:

$$R(\omega) = \int L(y, f(\mathbf{x}, \omega)) p(\mathbf{x}, y) dx dy \quad (2.7)$$

Learning is the process of estimating the function $f(\mathbf{x}, \omega_0)$, which minimizes the risk functional over the set of functions supported by the learning machine using only the training data ($p(\mathbf{x}, y)$ is not known). With finite data we cannot expect to find $f(\mathbf{x}, \omega_0)$ exactly, so we denote $f(\mathbf{x}, \omega^*)$ as the estimate of the optimal solution obtained with finite training data using some learning procedure. It is clear that any learning task (regression, classification, etc.) can be solved by minimizing (2.7) if the density $p(\mathbf{x}, y)$ is known. This means that density estimation is the most general (and hence most difficult) type of learning problem. The problem of learning (estimation) from finite data alone is inherently ill-posed. To obtain a useful (unique) solution, the learning process needs to incorporate a priori knowledge in addition to data. Let us assume that a priori knowledge is reflected in the set of approximating functions of a learning machine (as discussed earlier in this section). Then the next issue is: How should a learning machine use training data? The answer is given by the concept known as an *inductive principle*. An inductive principle is a general prescription for obtaining an estimate $f(\mathbf{x}, \omega^*)$ of the "true dependency" in the class of approximating functions, from the available (finite) training data. An inductive principle tells us *what* to do with the data, whereas the learning method specifies *how* to obtain an estimate. Hence a learning method (or algorithm) is a constructive implementation of an inductive principle for selecting an estimate $f(\mathbf{x}, \omega^*)$ from a particular set of functions $f(\mathbf{x}, \omega)$. For a given inductive principle there are many learning methods corresponding to a different set of functions of a learning machine. The distinction between inductive principles and learning methods is further discussed in Section 2.3.

2.1.2 Common Learning Tasks

The generic learning problem can be subdivided into four classes of common problems: classification, regression, density estimation, and clustering/vector quantization. For each of these problems, the nature of the loss function and

EXHIBIT 10

Webster's
Third
New International
Dictionary
OF THE ENGLISH LANGUAGE
UNABRIDGED

© Merriam-Webster
REG. U. S. PAT. OFF.

*Utilizing all the experience and resources of more than
one hundred years of Merriam-Webster® dictionaries*

EDITOR IN CHIEF
PHILIP BABCOCK GOVE, Ph. D.
AND
THE MERRIAM-WEBSTER
EDITORIAL STAFF



MERRIAM-WEBSTER INC., *Publishers*
SPRINGFIELD, MASSACHUSETTS, U.S.A.



A GENUINE MERRIAM-WEBSTER

The name *Webster* alone is no guarantee of excellence. It is used by a number of publishers and may serve mainly to mislead an unwary buyer.

Merriam-Webster[™] is the name you should look for when you consider the purchase of dictionaries or other fine reference books. It carries the reputation of a company that has been publishing since 1831 and is your assurance of quality and authority.

COPYRIGHT © 2002 BY MERRIAM-WEBSTER, INCORPORATED

WEBSTER'S THIRD NEW INTERNATIONAL DICTIONARY
PRINCIPAL COPYRIGHT 1961

Library of Congress Cataloging in Publication Data
Main entry under title:

Webster's third new international dictionary of the English language,
unabridged: a Merriam-Webster/editor in chief, Philip Babcock Gove
and the Merriam-Webster editorial staff.

p. cm.

ISBN-13: 978-0-87779-201-7 (buckram)

ISBN-10: 0-87779-201-1 (buckram)

1. English language—Dictionaries. I. Gove, Philip Babcock,
1902–1972. II. Merriam-Webster, Inc.

PE1625.W36 1993

423-dc20

All rights reserved. No part of this book covered by the copyrights hereon may be reproduced or copied in any form or by any means—graphic, electronic, or mechanical, including photocopying, taping, or information storage and retrieval systems—without written permission of the publisher.

MADE IN THE UNITED STATES OF AMERICA

5455QKY0605

revenue or involving a governmental expenditure or government property - compare PRIVATE CALENDAR
union card n 1: a card certifying personal membership in good standing in a labor union 2: something felt to resemble a union card esp. in being a prerequisite to employment or in providing evidence of in-group status (the Ph.D. ... a union card for the teaching profession - Douglas Bush)
union catalog n: a library catalog combining in one series and usu. alphabetically by authors a number of catalogs or the contents of more than one library
union day n, usu cap U&D: the anniversary of the founding of the Union of Africa on May 31, 1910 observed in the Union as a legal holiday
union depot n: UNION STATION
union district n: a school district made by uniting two or more elementary or secondary school districts
union down adv: with the flag reversed so that its union is down (a flag flown union down is a signal of distress at sea)

uni-on-ear \yünyöni(ä)r, -iä n-s: a member or advocate of a union; esp.: a labor union executive
union elbow n: an elbow pipe union
uni-on-oid \yüniöid, -id\ adj [NL Unionidae]: of or relating to the Unionidae
2unioid \yüniöid n-s: a mollusk of the family Unionidae
uni-on-i-dae \yüniöidä, -ä n pl, cap [NL, fr. Union, Unio, type genus + -idae]: a very large family of freshwater mussels (suborder Submytilacea) having a nearly often roughly sculptured shell with a thick epidermis and larvae that pass through a glochidium stage and being represented in nearly all parts of the world but chiefly in No. America where the naresous shells of many of them are used for button making
unionides pl of UNIO

uni-on-ism \yünyöniözəm n-s: the principle or policy of forming or adhering to a union: an advocacy or movement in favor of union; as: a usu cap: adherence to the policy of a firm federal union between the states of the United States esp. during the Civil War period b: the principles, theory, or system of combination of workers in the same occupation, trade, or industry (horizontal ~); also: the labor union movement (the advance of ~) c: advocacy of the principles of the British Unionists
uni-on-ist \yüniöist n: an advocate or promoter of union and esp. of some form of unionism; as: a usu cap: one loyal to the federal union of the U.S. during the Civil War b: a member of a former British political party advocating legislative union between Great Britain and Ireland c: an adherent or supporter of the labor union movement; esp.: an active member of a labor union d: an advocate of religious union and esp. of the amalgamation of related Protestant sects
uni-on-is-tic \yüniöistik, -ik\ adj: of, relating to, characteristic of, or favoring union or unionists

uni-on-ization \yünyöniözäshən, -äniöz- n-s 1: the quality or state of being unionized 2: the act of unionizing
uni-on-ize \yüniöz, -niz, -iz, -iz\ vt -ed/-ing-s see -ize in Explan Notes [Union + -ize]: to cause to become a member of or subject to the rules of a labor union (planned to ~ the shop): form into a labor union (unionizing previously unorganized groups)
union jack n, often cap U&A: a jack consisting of the union of a national flag
union jet burner n: a gas burner in which two jets unite to produce a single flat flame
union joint n: a joint (as between pipes) formed by means of a union

union label n: an identifying mark attached to goods indicating that they have been produced by union labor or that particular goods or services have been sold or done by that labor
union list n: a usu. alphabetical catalog of periodicals or other serials that provides bibliographical information and locates files in libraries
union-made \yüniöid\ adj: made by union labor
unions pl of UNION

union school n: an elementary, secondary, or combined elementary and secondary school that serves a union district - compare CONSOLIDATED SCHOOL
union security n: guaranteed preservation of union status or revenues obtained through clauses in a labor contract and esp. through provisions for closed or union shop, maintenance of membership, preferential hiring, or the checkoff
union service n: a worship service sponsored jointly by two or more religious organizations or communities (urged to attend the union service in the Congregational or Baptist Churches)
union shona n, usu cap U&S: SHONA 3
union shop n: an establishment in which the employer by agreement is free to hire nonmembers as well as members of the union but retains nonmembers on the payroll only on condition of their becoming members of the union within a specified time (a union shop clause in a collective agreement) - compare CLOSED SHOP
union station n: a station used jointly by two or more railroad or other transport companies (as bus or truck lines)
union suit n: an undergarment with shirt and drawers in one piece

union tannage n: tannage by means of a mixture of vegetable tanning materials
union tee n: a T pipe fitting with a male or female union at one end of the main run
unios pl of UNIO
uni-oval \yüniöval, -väl or uni-ovular \yüniöval, -väl\ adj [uni- + oval or ovular]: MONOVULAR
uni-ovulate \yüniövulat, -vüt\ adj [uni- + ovulate]: having a single ovule or ovum
uni-p-a-ri \yüniöpäriä n-s [NL, fr. uni- + -pariä]: a woman who has borne one child
uni-parental \yüniöpäntäl\ adj [uni- + parental]: having or involving a single parent; esp.: PARTHENOGENETIC - uni-parentally \yüniöpäntäl\ adv

uni-p-a-rous \yüniöpäros\ adj [uni- + -parous] 1 a: producing but one egg or offspring at a time b: having produced but one offspring at a once (terrestrial mammals) 2: producing but one axis at each branching (as a cyme)
uni-partite \yüniöpärti\ adj [uni- + partite] 1: not divided or divisible into parts
uni-ped \yüniöpäd\ n-s [uni- + -ped]: one having only one foot or leg
uni-personal \yüniöpönsäl\ adj [uni- + personal]: existing as one person
uni-personalist \yüniöpösäl n: one who believes that the deity is unipersonal
uni-personality \yüniöpösäl n: the quality or state of being unipersonal
uni-phase \yüniöpäz\ adj [uni- + phase]: having but one phase (a ~ conflict); esp.: SINGLE-PHASE
uni-planar \yüniöpänär\ adj [uni- + planar]: lying or occurring in one plane - compare PLANAR
uniplanar motion n: motion of a rigid body or fluid such that each point or particle moves in a plane parallel to a given plane - called also two-dimensional motion
uni-pod \yüniöpöd\ n-s [uni- + -pod (as in tripod)]: a one-legged support (as for a camera)
uni-polar \yüniöpölär\ adj [uni- + polar]: having or oriented in respect to a single pole; as: a: having, produced by, or acting by a single magnetic or electrical pole b: of a nerve cell: having but one process (as a bipolar cell) c: based on or controlled by a single compelling factor (as ~ coalition in politics) - uni-polarity \yüniöpöläriti\ n



uni-pulse \yüniöpüls\ n [uni- + pulse]: a single wave: PULSE
uni-quantie \yüniöpantik\ adj [uni- + quantie]: of, relating to, or giving rise to a single quantum of energy
uni-quantum \yüniöpantik\ n [uni- + -quantum (as in quantum)]: 1 a: being the only one; sole (earning money whose object could be nothing but Cyril's welfare - Arnold Bennett) (has thus preserved the original and often ~ records - C. B. Parks) (you are a miracle, a wonder, a mystery ~ one single ~ and inimitable living thing - J.C. Powys) b: of a book: known to exist in no other copy 2: being without a like or equal: single in kind or excellence: UNEQUALLED (they stand alone ~ of supreme interest - A.B. Osborne) (as historian he knows that events, like persons, are ~ - M. Barzun) (remains singularly himself, a ~ lyricalist of the first water - I.L. Solomon) (an almost ~ experience - Havelock Ellis) (tendencies present in our contemporary world which make our own times somewhat ~ - M.B. Smith) (story of his life is considerably more ~ than most autobiographies - Dorothy C. Fisher) (the more we study him, the less ~ he seems - Harry Levin) (sometimes used with to (the problem of what to do with surplus women is by no means ~ to our own society - Ralph Linton) or with (by no means ~ with the song sparrow - Nature Mag.) 3: UNUSUAL, NOTABLE (possessed ~ ability in the raising of funds - C.F. Thwing) (the wife of a career diplomat has ~ opportunity to observe the world political scene - Ray Pierce) (a frankness ~ in literature - David James) (peace and privacy - R.W. Hatch) (cheap, nourishing, and ~ dining experience - H. Fielding) (the most ~ characteristic of that environment - R.A. Billington) (she's the most ~ person I ever met - Arthur Miller) (the most ~ theater in town - advt) 4: capable of being performed in only one way (the factorization of a number into its prime factors is ~) SYN see SINGLE, STRANGE

unique \yüniök\ n-s 1: something (as a specimen, thing, circumstance, or person) that is unique (the only one of its kind (mistaking the ~ of the type - Roy Bedichek) (a display of collector for possession of a ~ - Roy Bedichek) (a display of glass, including undercooked ~s - Danish Foreign Office Jour.) (the phoenix, the ~ of birds - Thomas De Quincey)
unique-ly \yüniökli\ adv: in a unique manner 2: so as to be unique
unique-ness \yüniöknes\ n: the quality or state of being unique
unique-ness theorem n: a theorem in mathematics: a given point is a singular point on a surface
uni-que-ly \yüniökli\ adv: in a unique manner 2: so as to be unique
uni-radiate \yüniöräd-iät\ adj [uni- + radiate]: having a single ray or radius; esp.: MONAXON
uni-ramous \yüniörä-müs\ or uni-ramose \yüniörä-müz\ adj [uni- + ramous or ramose]: consisting of a simple process: UNBRANCHED (the appendages of crustaceans may be ~ or biramous)
uni-ramous \yüniörä-müs\ n-s [uni- + -ramous (as in trirame)]: a galley having but one tier of oars
un-ironed \yüniöränäd\ adj [un- + ironed, past part. of iron] 1: not restrained or confined with fetters 2: not pressed with a flatiron
unis abbr union

uni-serial \yüniörisäl\ or uni-seriate \yüniöris-iät\ adj [uni- + serial or seriate]: forming or arranged in a single series: having parts in a single row or on one side of an axis
uni-sexual \yüniösökshäl\ adj [uni- + -sexual]: of, relating to, or restricted to one sex: a: male or female but not hermaphroditic b: DICHLOUS (as ~ flower) - uni-sexuality \yüniösökshäl\ n - uni-sexually \yüniösökshäl\ adv
uni-son \yüniösön\ n-s [MI, fr. ML unisonus, adj., having the same sound, fr. L uni- + sonus sound - more at sound] 1: archaic: a tone identical in pitch with another; also: a whole or part of a scale 2: a: a pitch 2: a: a pitch identical in pitch with another; also: a whole or part of a scale 3: the interval of a perfect prime represented by the vibration ratio of 1:1 b: the state of being so tuned or sounded c: the writing, playing, or singing of parts in a musical passage at the same pitch or in octaves - compare HARMONY 2b 3 a: a harmonious agreement or union: CONCORD b: an instance or means of such agreement: a sympathy or sympathy: ASSENT (-s of overmastering thoughts) - in unison adv: in precise or perfect agreement: so as to harmonize exactly (speaking in unison)
unison \yüniösön\ n: 1: archaic: a: CONCORDANT, CONSONANT b: EQUIVALENT 2 a: identical in musical pitch: UNISONOUS (-singing) (as ~ passage) b: tuned to the same pitch - used of a string and esp. of any one of two or three piano strings that are struck by one hammer (~ strings) c: having a pitch that corresponds with the notation (as of a pipe-organ stop)
unis-on-al \yüniösönäl\ adj [unison + -al]: UNISONOUS
unis-ö-ö-nance \yüniösönäns\ n-s [unison + -ance]: a blending of sound into unison (the first two sounds reach the ear as a ~ - Newsweek)
unis-ö-ö-nant \yüniösönänt\ adj [unison + -ant]: UNISONOUS 1
unis-ö-ö-nous \yüniösönöüs\ n-s [MI, fr. ML unisonus] 1: in unison - used as a direction esp. in ensemble instrumental music

unis-ö-ö-nous \yüniösönöüs\ adj [ML unisonus] 1: being in unison: having the same degree of gravity or acuteness; sounded alike in pitch 2: alike in nature: CONCORDANT
un-issued stock \yüniöüsüd\ n [un- + issued, past part. of issue]: stock authorized (as under the charter of a corporation) but not yet issued - compare TREASURY STOCK
unit \yüniät, -üt, -äd, -ä n n-s [back-formation fr. unity] 1 (1) the smallest number or number that is the least whole number and is expressed by the numeral 1 (2) a single thing (as a magnitude or number) that constitutes an undivided whole b: a number that divides every element of a set of numbers c: a determinate quantity (as of length, time, heat, value, or housing) adopted as a standard of measurement for other quantities of the same kind: as (1) a fractional part of the width of a printing character (as 1/4 of a fractional roman capital M) used in measuring the set of a piece of type and being of the same width for all type of the same point size and proportionally wider or narrower for larger or smaller point sizes (2) an amount of work (as 120 hours of classroom work in a completed course of a secondary school) used in education in calculating student credits (as for graduation or college entrance) (3) an amount of a biologically active agent (as a vitamin, or antigen) required to produce a specific result under strict controlled conditions compare BIOASSAY, RAT UNIT (4) one percent per ton of a fertilizing ingredient (a fertilizer containing 5 percent of nitrogen, 10 percent of phosphoric acid, and 10 percent of potash includes 25 fertilizer ~s) 2 a: a single thing or person or group that is a constituent and isolable member of some more inclusive whole b: a member of an aggregate that is the least clearly definable separate existence and that normally forms a basic element of organization within the aggregate (the township in the usual ~ of government) (the family as a basic ~ of society) b: one of the commonly more or less repetitive sections combined in assembling a manufactured article (as a bookcase or kitchen cabinet) c: a part of a military establishment that has a prescribed organization (as of personnel and materiel) (in the army ~s vary in size and complexity from the squad to the army) d: a piece or complex of apparatus serving a single or common purpose (as an oil train drawn by two diesel ~s) (a power station with one ~ out of order) e: a combination of two or more securities offered at a single price (as ~ of one share of preferred stock and two shares of common offered at \$110) f: a course or part of a course in an elementary or secondary school focusing on a central theme and making use of resources from numerous subjects (as the pupils' own experience g: BARGAINING UNIT h: a fraction of an annual pension or a retirement income benefit earned as a result of each year's service prior to

unit-able also unite-able \yüniät-äbel, -Itä- adj [unite + -able]: capable of union by growth or otherwise: that can be joined together
unit-age \yüniät-äje n-s [unit + -age] 1: specification of the amount constituting a unit (as of a vitamin) 2: amount in units (as ~ of 50,000 per capsule)
unit-al \yüniätäl\ adj [unit + -al]: UNITARY
uni-tar-i-an \yüniätäriän, -tä(ä)r- n-s [NL unitarius]
unitarian (fr. L unitas - past part. of unire to unite - + -arius -ary) + E-on 1: a usu cap a: a Christian who believes that the deity exists only in one person: a unipersonalist who denies the doctrine of the Trinity; a member of a Christian denomination who in general affirms the principles of individual freedom of belief, freedom of religion, religious tolerance, commitment to man, a creedless church, a united world community, and support of a vigorous program of liberal social action b: a non-Christian monotheist (as a Muhammadan) 2 a: an advocate of a theory or doctrine founded upon unity; specif.: MONIST b: an advocate of unity or a unitary system; specif.: one who advocates centralization in government c: an observer of the

unitarian \yüniätäriän\ adj 1: a usu cap: of, relating to, or involving Unitarians or their doctrines 2: UNITARY 1: as a MONIST b: of, relating to, or advocating centralization in government or administration
unitarian hypothesis also unitarian view n: a theory in immunology: a single pure antigen will produce only one variety of antibody, which when brought into contact with the antigen in appropriate form can react in various ways (as by agglutinating, precipitating, fixing complement, or opsonizing)
uni-tar-i-an-ism \yüniätäriänizim\ n-s 1 cap: the principles and practices of Unitarians 2 sometimes cap: a unitarian or unitary system (as of government)
uni-tar-i-ness \yüniätäriännes, -riti- n-ES: the quality or state of being unitary
uni-tar-y \yüniätäri\ adj [unit & unity + -ary] 1: of, relating to, based upon, or characterized by unity (as ~ movement in politics) 2: MONISTIC 2 a: having the character of a unit: not divided or discontinuous (as ~ process) b: functioning as a unit esp. of measurement (established a ~ distance on which to base subsequent calculations) (a ~ university) c: of, relating to, or constituting a system of government in which power is held by a central authority and may be delegated, but not derived from constituent subdivisions - extended: of, relating to, or involving the use of units (as ~ approach to a problem) (the methods ~ arithmetic) b: made up of discrete units (as ~ loudspeaker with four speaker units associated in a single assembly)
unitary color n: PSYCHOLOGICAL PRIMARY
unitary theory n: a theory in chemistry: molecules are units whose parts are bound together in definite structure with mutual and reciprocal influence on each other - compare DUALISM

unit banking n: banking carried on by individual banks without branches or corporate relationships with other banks
unit card n: a library catalog card containing full information about a book or other printed item and reproduced in quantity so that it may be not only used for a main entry but adapted for all secondary entries
unit cell n: the smallest polyhedron that by indefinite repetition makes up the lattice of a crystal and embodies all the characteristics of its structure
unit character n 1: a natural character that is inherited on an all or none basis 2: a natural character dependent on the presence or absence of a single gene: a typical Mendelian or qualitative character
unit class n: a class with a single member
unit construction n: a system of building in which large sections (as of a ship) can be fabricated independently and subsequently assembled
unit cost n: the cost allocated to a selected unit and commonly calculated as the cost over a period of time divided by the number of items produced
unite \yüniät, -üt, -äd, -ä vb -ed/-ing-s [ME unire, fr. LL unius, past part. of unire, fr. L unis one, unire to unite, fr. L un- + -ire: to put together a single unit into the fighting forces of the friendly nations] b: to cause to adhere (~ bricks with mortar) c: CONNECT (a dirt road ~s the farm road with the main highway) d: to relate integrally (often the ideas are yoked, but not united - T.S. Eliot); esp.: to link by a legal or moral bond (a purpose that united all factions) (a treaty to ~ all the independent nations) 2: to possess (as qualities) in combination (the bride united beauty and intelligence) ~ vt 1 a: to become one or as if one (particles which can ~ to form a new compound - I.S. Eliot) (mutterings of the crowd united in a thunderous cheer - Darrell Berrigan) b: to become combined by or as if by adhesion or mixture (the broken bones of a child ~ easily) (clouds of devastating smoke that ~ with the river fog ~ to form smog - Amer. Guide Series: Pa.) 2: to act in concert (all parties united in signing the petition) 3: to enter into a common endeavor or as if for a common purpose (the group united to improve the city's schools)

syn COMBINE, CONJOIN, CONCUR, COOPERATE: UNITE often indicates joining, merging, coalescing, adhering together to form a new unit, permanent or temporary (the North West Company united with the Hudson's Bay Company - Amer. Guide Series: Wash.) (in France the whole people was united to unite their country in a common endeavor - Hilaire Belloc) W.R. Inge) COMBINE may apply to a temporary uniting or joining or to one which leaves the components distinct (a gift for combining, for fusing into a single phrase, two or more diverse impressions - T.S. Eliot) (wealth and sophistication combine with breezy western characteristics in this town - Amer. Guide Series: Texas) (innumerable factors combine in the intricate complexity of our general story - Hilaire Belloc) CONJOIN is like UNITE in stressing the notion of jointure, often of more or less equal things or forces, at a specific point (nature had lavished gifts and aspirations upon him, but they were so mixed and contradictory that only by a fortunate miracle did he become one of them conjoined to produce the rich poetry by which he is remembered - R.D. Altick) CONCUR is likely to coincide when things that happen to merge, work together, or are of one order (the conflict in a common endeavor or plausible (two opposite forces concur in bringing about the Council of Nicea - A.P. Stanley) COOPERATE indicates a joining of strength or force in some specific situation with no fusion or loss of identity (sent a joint expedition, under British command, to cooperate with the White Russians at Murmansk and Archangel against the Bolshevik forces - J.M. Hanson) SYN see in addition JOIN

unite \yüniät, -üt, -äd, -ä n-s [fr. obs. unium, united, fr. ME unit, fr. L unius] an old British gold 20-shilling piece issued first by James I in 1604 for England and Scotland and bearing in the design and inscription reference to the uniting of the two crowns - called also JACOBUS
unite-able also unite-able \yüniät-äbel, -Itä- adj [unite + -able]: capable of union by growth or otherwise: that can be joined together
unit-age \yüniät-äje n-s [unit + -age] 1: specification of the amount constituting a unit (as of a vitamin) 2: amount in units (as ~ of 50,000 per capsule)
unit-al \yüniätäl\ adj [unit + -al]: UNITARY
uni-tar-i-an \yüniätäriän, -tä(ä)r- n-s [NL unitarius]
unitarian (fr. L unitas - past part. of unire to unite - + -arius -ary) + E-on 1: a usu cap a: a Christian who believes that the deity exists only in one person: a unipersonalist who denies the doctrine of the Trinity; a member of a Christian denomination who in general affirms the principles of individual freedom of belief, freedom of religion, religious tolerance, commitment to man, a creedless church, a united world community, and support of a vigorous program of liberal social action b: a non-Christian monotheist (as a Muhammadan) 2 a: an advocate of a theory or doctrine founded upon unity; specif.: MONIST b: an advocate of unity or a unitary system; specif.: one who advocates centralization in government c: an observer of the

EXHIBIT 11

The Oxford English Reference Dictionary

Second Edition

Edited by

Judy Pearsall and Bill Trumble

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide in

Oxford New York

Athens Auckland Bangkok Bogota Buenos Aires Calcutta
Cape Town Chennai Dar es Salaam Delhi Florence Hong Kong Istanbul
Karachi Kuala Lumpur Madrid Melbourne Mexico City Mumbai
Nairobi Paris São Paulo Singapore Taipei Tokyo Toronto Warsaw
with associated companies in Berlin Ibadan

Oxford is a registered trade mark of Oxford University Press
in the UK and in certain other countries

© Oxford University Press 1995, 1996

First edition 1995

Second edition 1996

All rights reserved. No part of this publication may be reproduced,
stored in a retrieval system, or transmitted, in any form or by any means,
without the prior permission in writing of Oxford University Press,
or as expressly permitted by law, or under terms agreed with the appropriate
reprographics rights organization. Enquiries concerning reproduction
outside the scope of the above should be sent to the Rights Department,
Oxford University Press, at the address above

You must not circulate this book in any other binding or cover
and you must impose this same condition on any acquirer

British Library Cataloguing in Publication Data

Data available

Library of Congress Cataloging in Publication Data

Data available

ISBN 0-19-860046-1 Plain Edition

ISBN 0-19-860050-X Thumb Index Edition

10 9 8 7 6 5

Printed on acid-free paper in the USA

of a person or thing. **4** (of a disease) with no known external stimulus or cause; idiopathic. **n.** (esp. in *pl.*) a basic or indispensable element or thing. **□ essential element** a chemical element required by living organisms for normal growth. **essential oil** an oil present in and having the characteristic odour of a plant etc., from which it can be obtained by distillation (see also *oil*). **□ essentially adv.** **essentialness n.** **essentiality** /iːsenʃiəlɪti/ *n.* [ME *f.* *LL. essentialis* (as *ESSENCE*)]

Essequibo /esɪˈkiːbəʊ/ a river in Guyana, rising in the Guiana Highlands and flowing about 965 km (600 miles) northwards to the Atlantic.

Essex /'esɪks/ a county of eastern England; county town, Chelmsford.

EST *abbr.* 1 Eastern Standard Time. 2 electro-shock treatment.

-est /ɪst/ *suffix* forming the superlative of adjectives (*widest; nicest; happiest*) and adverbs (*soonest*). [OE *-ost-, -ust-, -ast-*]

-est /ɪst/ *suffix* (also **-st**) *archaic* forming the 2nd person sing. of verbs (*canst; findest; gavest*). [OE *-est-, -ast-, -st-*]

establish /ɪ'stæblɪʃ/ *vtr.* 1 set up or consolidate (a business, system, etc.) on a permanent basis. 2 (foll. by *in*) settle (a person or oneself) in some capacity. 3 (esp. as **established adj.**) achieve permanent acceptance for (a custom, belief, practice, institution, etc.). 4 **a** validate; place beyond dispute (a fact etc.). **b** find out, ascertain. **□ established Church** a Church recognized by the state as the national Church. **□ establisher n.** [ME *f.* OF *establiŕ* (stem *establiŕs*) *f.* *L. stabiliŕe f. stabiliŕs* STABLE]

establishment /ɪ'stæblɪʃmənt/ *n.* 1 the act or an instance of establishing; the process of being established. 2 **a** a business organization or public institution: **b** a place of business. **c** a residence. 3 **a** the staff or equipment of an organization. **b** a household. 4 any organized body permanently maintained for a purpose. 5 a Church system organized by law. 6 **a** (the Establishment) the group in a society exercising authority or influence, and seen as resisting change. **b** any influential or controlling group (the literary Establishment).

establishmentarian /ɪ'stæblɪʃməntɪ'æəriən/ *adj. & n.* **• adj.** adhering to or advocating the principle of an established Church. **• n.** a person adhering to or advocating this. **□ establishmentarianism n.**

estaminet /c'estæmiːneɪ/ *n.* a small French café etc. selling alcoholic drinks. [F *f.* Walloon *staminé* *byre f. stamo* a pole for tethering a cow, *prob. f. G Stamm stem*]

estate /ɪ'steɪt/ *n.* 1 a property consisting of an extensive area of land usu. with a large house. 2 *Brit.* a modern residential or industrial area with integrated design or purpose. 3 all of a person's assets and liabilities, esp. at death. 4 a property where rubber, tea, grapes, etc., are cultivated. 5 (in full **estate of the realm**) an order or class forming (or regarded as) a part of the body politic. 6 *archaic* or *literary* a state or position in life (the *estate of holy matrimony; poor man's estate*). 7 *colloq.* = *estate car*. **□ estate agent** *Brit.* 1 a person whose business is the sale or lease of buildings and land on behalf of others. 2 the steward of an estate. **estate car** *Brit.* a car with the passenger area extended and combined with space for luggage, usu. with an extra door at the rear. **estate duty** *Brit. hist.* death duty levied on property. ¶ Replaced in 1975 by *capital transfer tax* and in 1986 by *inheritance tax*. [ME *f.* OF *estat* (as *STATUS*)]

Estates General see STATES GENERAL.

esteem /ɪ'stiːm/ *v. & n.* **• vtr.** 1 (usu. in *passive*) have a high regard for; greatly respect; think favourably of. 2 *formal* consider, deem (esteemed it an honour). **• n.** high regard; respect; favour (*held them in esteem*). [ME *f.* OF *estimer f. L aestimare* fix the price of]

ester /'estə(r)/ *n.* *Chem.* an organic compound produced by replacing the hydrogen of an acid by an alkyl, aryl, etc., radical, many examples of which occur naturally as oils and fats. **□ esterify** /c'estəriːfaɪ/ *vtr.* (-ies, -ied). [G, *prob. f. Essig vinegar + Ather ether*]

Esth. *abbr.* (in the Bible & Apocrypha) Esther.

Esther /'estə(r)/ 1 (in the Bible) a woman who was chosen on account of her beauty by the Persian King Ahasuerus (generally supposed to be Xerxes I) to be his queen and who used her influence with him to save the Israelites in captivity from persecution. 2 the book of the Bible containing an account of these events; a part survives only in Greek and is included in the Apocrypha.

esthete *US var.* of AESTHETE.

aesthetic *US var.* of AESTHETIC.

estimable /'estɪməb(ə)l/ *adj.* worthy of esteem. **□ estimably adv.** [F *f.* *L aestimabilis* (as *ESTEME*)]

estimate *n. & v.* **• n.** /'estɪmət/ 1 an approximate judgement, esp. of cost, value, size, etc. 2 a price specified as that likely to be charged for work to be undertaken. 3 opinion, judgement, estimation. **• vtr.** (also *absol.*) /'estɪmeɪt/ 1 form an estimate or opinion of. 2 (foll. by *that* + *clause*) make a rough calculation. 3 (often foll. by *a*) value or measure by estimation: *adjudge*. **□ estimative** /-mətɪv/ *adj.* **estimator** /-meɪtə(r)/ *n.* [*L aestimare aestimat*: fix the price of]

estimation /'estɪmeɪʃ(ə)n/ *n.* 1 the process or result of estimating. 2 judgement or opinion of worth (in *my estimation*). 3 *archaic* esteem (hold in estimation). [ME *f.* OF *estimation* or *L aestimatio* (as *ESTIMATE*)]

estivate *US var.* of AESTIVAL.

estivate *US var.* of AESTIVATE.

Estonia /ɪ'stəʊniə/ a Baltic country on the south coast of the Gulf of Finland; pop. (est. 1991) 1,591,000; languages, Estonian (official), Russian; capital, Tallinn. Estonia is a flat, lowland country with marshland, lakes, and forest. Previously ruled by the Teutonic Knights and then by Sweden, Estonia was ceded to Russia in 1721. It was proclaimed an independent republic in 1918 but was annexed by the USSR in 1940 as a constituent republic, the Estonian SSR. With the breakup of the Soviet Union Estonia regained its independence in 1991.

Estonian /ɪ'stəʊniən/ *n. & adj.* **• n.** 1 **a** a native of Estonia. **b** a person of Estonian descent. 2 the Finno-Ugric language of Estonia, most closely related to Finnish and spoken by about a million people. **• adj.** of or relating to Estonia or its people or language.

estop /ɪ'stɒp/ *vtr.* (**estopped, estopping**) (foll. by *from*) *Law* bar or preclude, esp. by estoppel. **□ estoppage n.** [ME *f.* AF, OF *estoper f. LL stuppare* stop up *f. L stuppa* tow: cf. *strop, struff*]

estoppel /ɪ'stɒp(ə)l/ *n.* *Law* the principle which precludes a person from asserting something contrary to what is implied by a previous action or statement of that person, or by a previous pertinent judicial determination. [OF *estouppail* bung *f. estoper* (as *ASTOP*)]

Estoril /'estɔːrɪl/ a resort on the Atlantic coast of Portugal; pop. (1991) 24,850.

estovers /ɪ'stəʊvəz/ *n. pl. hist.* necessities allowed by law to a tenant (esp. fuel, or wood for repairs). [AF *estover*, OF *estoveir* be necessary, *f. L est opus*]

estrangle /ɪ'streɪŋdʒ/ *vtr.* (usu. in *passive*; often foll. by *from*) 1 cause (a person or group) to turn away in feeling or affection; alienate. 2 (as **estranged adj.**) (of a husband or wife) no longer living with his or her spouse. **□ estrangement n.** [ME *f.* AF *estranger*, OF *estranger f. L extraneare* treat as a stranger *f. extraneus* stranger]

estreat /ɪ'stri:t/ *n. & v.* *Law* **• n.** 1 a copy of a court record of a fine etc. for use in prosecution. 2 the enforcement of a fine or forfeiture of a recognizance. **• vtr.** enforce the forfeit of (a fine etc., esp. surety for bail). [ME *f.* AF *estrete*, OF *estraite f. estraire f. L extrahere* EXTRACT]

Estremadura /'estɹəmə'duərə/ a coastal region and former province of west central Portugal.

estrogen *US var.* of OESTROGEN.

estrus etc. *US var.* of OESTRUS etc.

estuary /'estʃuəri/ *n.* (*pl.* -ies) a wide tidal mouth of a river. **□ estuarine** /-rɪn/ *adj.* [*L aestuarium* tidal channel *f. aestus* tide]

e.s.u. *abbr.* electrostatic unit(s).

esurient /ɪ'sjʊəriənt/ *adj. archaic* or *joc.* 1 hungry. 2 impecunious and greedy. **□ esuriently adv.** [*L esurire* (v.) hunger *f. edere* eat]

Esztergom /'estɔːɡɒm/ a town and river port on the Danube in Hungary; pop. (est. 1984) 31,000.

ET *abbr.* extraterrestrial.

-et /ɪt/ *suffix* forming nouns (orig. diminutives) (*baronet; bullet; sonnet*). [OF *-ete*]

-et /ɪt/ *suffix* (also **-ete** /i:t/) forming nouns usu. denoting persons (*comet; poet; athlete*). [Gk *-ētes*]

ETA *abbr.* estimated time of arrival.

ETA /'etə/ a Basque separatist movement in Spain which has waged a terrorist campaign since its foundation in 1959 for an independent Basque state. [Basque acronym, *f. Euzkadi ta Azkatasuna* Basque homeland and liberty]

eta /'i:tə/ *n.* the seventh letter of the Greek alphabet (*Η, η*). [Gk]

et al. /et 'æl/ *abbr.* and others. [*L et alii, et alia*, etc.]

etalon /'etəˌlɒn/ *n.* *Physics* a device consisting of two reflecting plates for producing interfering light-beams. [*F étalon* standard]

etc. *abbr.* = ET CETERA.

et cetera /et 'setərə, 'setrə/ *rest*; and similar things or p so on. **• n.** (in *pl.*) the usual

etch /etʃ/ *v. & n.* **• v.** 1 **a** *tr.* design on a metal plate wit plate) in this way. 2 *intr.* prac deeply (esp. on the mind); **□ etcher n.** [Du. *etsen f. G ä eaten f. Gmc*]

etchant /'etʃənt/ *n.* a corr

etching /'etʃɪŋ/ *n.* 1 a pri producing such plates. Th century, though the basi metal plate, had been used

-ete *suffix* var. of -ET.

eternal /ɪ'tɜːn(ə)l/ *adj.* 1 beginning in time. 2 esse constant; seeming not t

triangle a relationship **□ eternalize vtr.** (also

eternity /ɪ'tɜːnəlɪti/ *n.* [age]

Eternal, the God.

Eternal City, the Rome

eternity /ɪ'tɜːnɪti/ *n.* (*pl.* -I 2 (in Christian theology) eternal. 4 (often prec. by truths. **□ eternally ring**

given as a token of lasti -atis *f. aeternus*: see ETERN

Etesian winds /ɪ'teɪz summer in the eastern M etos year]

eth /eθ/ *n.* (also **edh** /eð

letter, ð, capital ð (= th)

-eth *var.* of -TH.

-eth /ɪθ/ *suffix* (also **-th**) *ε verbs* (*doeth; saith*). [OE *-ethanal* /'eθəˌnæl, 'i:θ-.

ethane /'i:θem, 'eθ-/ *n.* series (chem. formula: C

ethanediol /'i:θem,dar

ethanoic acid /'eθəˌn /-noʊeɪt/ *n.* [ETHANE + -I

ethanol /'eθəˌnɒl, 'i:θ-

Ethelred /'eθəl,rəd/ *th*

Ethelred II (known a good advice; rash) (c.96 inability to confront th brother St Edward the

their attacks. In 1013 Sweyn I.

ethene /'eθiːn, 'i:θ-/ *n.*

ether /'i:θə(r)/ *n.* 1 *Chem.* organic liquid (chem. solvent. Also called *eth*

with a similar structu etc. groups. 2 (also *aet*

the clouds. 3 (also *aet* permeate space and *et*

medium through whic to be transmitted. **□ e**

Gk either *f.* root of *aith*

etheral /'i:θərəl/ *ac* esp. in appearance. 3 **□ etherally adv.** *et*

aetherios (as *ETHER*)]

etherial *var.* of ETHER

etherize /'i:θəˌraɪz/ *v* **□ etherization** /i:θəˌraɪz/ *n.* **Ethernet** /'i:θəˌnet/ *n.* networks using coaxi

probabilistic /prɒbəbəlɪstɪk/ *adj.* relating to probability; involving chance variation.

probability /prɒbəbɪləti/ *n.* (pl. **-ies**) 1 the state or condition of being probable. 2 the likelihood of something happening. 3 a probable or most probable event (the probability is that they will come). 4 *Math.* the extent to which an event is likely to occur, measured by the ratio of the favourable cases to the whole number of cases possible. □ **In all probability** most probably. [F *probabilité* or L *probabilitas* (as PROBABLE)]

probable /ˈprɒbəb(ə)l/ *adj.* & *n.* ● *adj.* (often foll. by *that* + clause) that may be expected to happen or prove true; likely (the probable explanation; it is probable that they forgot). ● *n.* a probable candidate, member of a team, etc. □ **probably** *adv.* [ME f. OF f. L *probabilis* f. *probare* prove]

proband /ˈprɒbənd/ *n.* a person forming the starting-point for the genetic study of a family etc. [L *probandus*, gerundive of *probare* test]

probang /ˈprɒbəŋ/ *n.* *Surgery* a strip of flexible material with a sponge etc. at the end, used to remove a foreign body from the throat or apply a medication to it. [17th c. (named *provang* by its inventor): orig. unkn., *perh.* alt. after *probe*]

probate *n.* & *v.* ● *n.* /ˈprəʊbət, -bət/ 1 the official proving of a will. 2 a verified copy of a will with a certificate as handed to the executors. ● *vtr.* /ˈprəʊbət/ *N. Amer.* establish the validity of (a will). [ME f. L *probatum* neut. past part. of *probare* PROVE]

probation /prəˈbeɪʃ(ə)n/ *n.* 1 *Law* a system of supervising and monitoring the behaviour of (esp. young) offenders, as an alternative to prison. 2 a process or period of testing the character or abilities of a person in a certain role, esp. of a new employee. 3 a moral trial or discipline. □ **on probation** undergoing probation, esp. legal supervision. □ **probation officer** an official supervising offenders on probation. □ **probational** *adj.* **probationary** *adj.* [ME f. OF *probation* or L *probatio* (as PROVE)]

probationer /prəˈbeɪʃənə(r)/ *n.* 1 a person on probation, e.g. a newly appointed nurse, teacher, etc. 2 an offender on probation. □ **probationership** *n.*

probative /ˈprəʊbətɪv/ *adj.* affording proof; evidential. [L *probativus* (as PROVE)]

probe /prəʊb/ *n.* & *v.* ● *n.* 1 a penetrating investigation. 2 any small device, esp. an electrode, for measuring, testing, etc. 3 a blunt-ended surgical instrument usu. of metal for exploring a wound etc. 4 (in full **space probe**) an unmanned exploratory spacecraft transmitting information about its environment. ● *v.* 1 *tr.* examine or enquire into closely. 2 *tr.* explore (a wound or part of the body) with a probe. 3 *tr.* penetrate with or as with a sharp instrument, esp. in order to explore. 4 *intr.* make an investigation with or as with a probe (the detective probed into her past life). □ **probeable** *adj.* **prober** *n.* **probingly** *adv.* [LL *proba* proof, in med.L = examination, f. L *probare* test]

probit /ˈprɒbɪt/ *n.* *Statistics* a unit of probability based on deviation from the mean of a standard distribution. [probability unit].

probity /ˈprəʊbɪti/ *n.* uprightness, honesty. [F *probité* or L *probitas* f. *probus* good]

problem /ˈprɒbləm/ *n.* 1 a doubtful or difficult matter requiring a solution (how to prevent it is a problem; the problem of ventilation). 2 something hard to understand or accomplish or deal with. 3 (attrib.) causing problems; difficult to deal with (problem child). 4 *a Physics & Math.* an inquiry starting from given conditions to investigate or demonstrate a fact, result, or law (cf. THEOREM 1). 5 *Geom.* a proposition in which something has to be constructed. 5 *a* (in various games, esp. chess) an arrangement of men, cards, etc., in which the solver has to achieve a specified result. 5 *a* puzzle or question for solution. □ **that's your (or his etc.) problem** said to disclaim responsibility or connection. [ME f. OF *probleme* or L *problema* f. Gk *problēma* -matos f. *proballō* (as PRO-², ballō throw)]

problematic /prɒbləˈmætɪk/ *adj.* (also **problematical** /-k(ə)l/) 1 attended by difficulty. 2 doubtful or questionable. 3 *Logic* enunciating or supporting what is possible but not necessarily true. □ **problematically** *adv.* [F *problématique* or LL *problematicus* f. Gk *problēmatikos* (as PROBLEM)]

proboscidean /prəʊbəsɪdiən/ *adj.* & *n.* (also **proboscidian**) *Zool.* ● *adj.* 1 having a proboscis. 2 of or like a proboscis. 3 of or relating to the mammalian order Proboscidea, which includes elephants and related extinct animals. ● *n.* a mammal of this order. [mod.L *Proboscidea* (as PROBOSCIS)]

proboscis /prəʊˈbɒsɪs/ *n.* 1 the long flexible trunk or snout of some mammals, e.g. an elephant or tapir. 2 the elongated mouthparts of some insects, used for sucking liquids or piercing. 3 the sucking organ

in some worms. 4 *loc.* the human nose. □ **proboscis monkey** a monkey, *Nasalis larvatus*, native to Borneo, the male of which has a large pendulous nose. □ **probosciferous** /-bɒsɪˈdɪfərəs/ *adj.* **probosciform** /prəʊbəsɪˈfɔ:m/ *adj.* [L *proboscis*-cidis f. Gk *proboskīs* f. *proboskō* (as PRO-², boskō feed)]

procaine /ˈprəʊkeɪn/ *n.* (also **procain**) a synthetic compound used as a local anaesthetic, esp. in dentistry. [PRO-¹ + COCAINE]

procaryote var. of PROKARYOTE.

procedure /prəˈsi:dʒə(r), -sɪːdʒə(r)/ *n.* 1 a way of proceeding, esp. a mode of conducting business or a legal action. 2 a mode of performing a task. 3 a series of actions conducted in a certain order or manner. 4 a proceeding. 5 *Computing* = SUBROUTINE. □ **procedural** *adj.* **procedurally** *adv.* [F *procédure* (as PROCEED)]

proceed /prəˈsi:d/ *vitr.* 1 (often foll. by *to*) go forward or on further; make one's way. 2 (often foll. by *with*, or *to* + *inf.*) continue; go on with an activity (proceeded with their work; proceeded to tell the whole story). 3 (of an action) be carried on or continued (the case will now proceed). 4 adopt a course of action (how shall we proceed?). 5 go on to say. 6 (foll. by *against*) start a lawsuit (against a person). 7 (often foll. by *from*) come forth or originate (shouts proceeded from the bedroom). 8 (foll. by *to*) *Brit.* advance to a higher rank, university degree, etc. [ME f. OF *proceder* f. L *procedere* process- (as PRO-¹, *cedere* go)]

proceeding /prəˈsi:dɪŋ/ *n.* 1 an action or piece of conduct (a high-handed proceeding). 2 (in *pl.*) (in full **legal proceedings**) an action at law; a lawsuit. 3 (in *pl.*) a published report of discussions or a conference. 4 (in *pl.*) business, actions, or events in progress (the proceedings were enlivened by a dog running on to the pitch).

proceeds /ˈprəʊsi:dz/ *n.pl.* money produced by a transaction or other undertaking. [*pl.* of obs. *proceed* (n.) f. PROCEED]

process 1 /ˈprəʊses/ *n.* & *v.* ● *n.* 1 a course of action or procedure, esp. a series of stages in manufacture or some other operation. 2 the progress or course of something (in *process* of construction). 3 a natural or involuntary operation or series of changes (the process of growing old). 4 an action at law; a summons or writ. 5 *Anat., Zool., & Bot.* a natural appendage or outgrowth on an organism. ● *vtr.* 1 handle or deal with by a particular process. 2 treat (food, esp. to prevent decay) (processed cheese). 3 *Computing* operate on (data) by means of a program. □ **In process** going on, being done. **In process of time** as time goes on. **process server** a sheriff's officer who serves writs. □ **processable** *adj.* [ME f. OF *proces* f. L *processus* (as PROCEED)]

process 2 /ˈprəʊses/ *vitr.* walk in procession. [back-form. f. PROCESSION]

procession /prəˈseɪʃ(ə)n/ *n.* 1 a number of people or vehicles etc. moving forward, in orderly succession, esp. at a ceremony, demonstration, or festivity. 2 the movement of such a group (go in procession). 3 a regular succession of things; a sequence. 4 a race in which no competitor is able to overtake another. 5 (in Christian theology) the emanation of the Holy Spirit. □ **processionist** *n.* [ME f. OF f. L *processio*-onis (as PROCEED)]

processional /prəˈseɪʃ(ə)n(ə)l/ *adj.* & *n.* ● *adj.* 1 of or relating to processions. 2 used, carried, or sung in processions. ● *n.* *Ecccl.* an office-book of processional hymns etc. [med.L *processionalis* (adj.), -ale (n.) (as PROCESSION)]

processor /ˈprəʊsesə(r)/ *n.* a machine or device that processes things, esp.: 1 *Computing* = central processor. 2 = food processor.

procès-verbal /ˈprəʊsɜ:vˈbɑ:l/ *n.* (pl. **procès-verbaux** /-bɔ:/) a written report of proceedings; minutes. [F]

pro-choice /prəʊˈtʃɔɪs/ *adj.* & *n.* ● *adj.* advocating a woman's legal right to choose whether to have an abortion. ● *n.* a pro-choice policy.

prochronism /prəʊkɹɒnɪz(ə)m/ *n.* the action of referring an event etc. to an earlier date than the true one. [PRO-² + Gk *khronos* time]

proclaim /prəˈkleɪm/ *vtr.* 1 (often foll. by *that* + clause) announce or declare publicly or officially. 2 declare (a person) to be (a king, traitor, etc.). 3 reveal as being (an accent that proclaims you a Scot). □ **proclaimer** *n.* **proclamatory** /-ˈklæmətɔ:ri/ *adj.* **proclamation** /ˌprɒkləˈmeɪʃ(ə)n/ *n.* [ME *proclame* f. L *proclamare* cry out (as PRO-¹, CLAIM)]

proclitic /prəˈklɪtɪk/ *adj.* & *n.* *Gram.* ● *adj.* (of a monosyllable) closely attached in pronunciation to a following word and having itself no accent. ● *n.* such a word, e.g. *at* in *at in at home*. □ **proclitically** *adv.* [mod.L *procliticus* f. Gk *proklitō* lean forward, after LL *encliticus*: see ENCLITIC]

proclivity /prəˈklɪvɪti/ *n.* (pl. **-ies**) a tendency or inclination. [L *proclivitas* f. *proclivis* inclined (as PRO-¹, *clivus* slope)]

Procne /ˈprɒkni/ *Gk Mythol.* the sister of Philomel.

proconsul /prəʊˈkɒns(ə)l/ *n.* 1 *Rom. Hist.* a governor of a province, in

the l:
etc. 5
adj. p
for t)
Proc
Caes
(c.50:
are t
auth
but i
poli
mor
proc
reas
□ pr
/-na
proc
pro
exis
pro
n. [I
Pro
forc
Pro
lie-
off
[G]
pro
the
/p
pro
un
fun
a f
ad
/p
/th
pre
th
pre
gr
la
pr
ot
of
pr
□
as
oi
[a
[a
pr
(r
a
p
(i
[i
pr
c
f
p
t
c
p
[i
[i
p
[i

EXHIBIT 12

THE AMERICAN HERITAGE
DICTIONARY
OF THE ENGLISH LANGUAGE



Words that are believed to be registered trademarks have been checked with authoritative sources. No investigation has been made of common-law trademark rights in any word, because such investigation is impracticable. Words that are known to have current registrations are shown with an initial capital and are also identified as trademarks. The inclusion of any word in this Dictionary is not, however, an expression of the publishers' opinion as to whether or not it is subject to proprietary rights. Indeed, no definition in this Dictionary is to be regarded as affecting the validity of any trademark.

© 1969, 1970, 1971, 1973, 1975, 1976, 1978 by Houghton Mifflin Company
All correspondence and inquiries should be directed to
Dictionary Division, Houghton Mifflin Company
One Beacon Street, Boston, Massachusetts 02107

All rights reserved under Bern and Pan-American Copyright Conventions


ISBN: 0-395-20360-0 (new college edition; thumb-indexed)
0-395-20359-7 (new college edition; plain edges)
0-395-24575-3 (high-school edition)

Library of Congress Catalog Card Number 76-86995

Manufactured in the United States of America

Computer-composed by Inforonics, Inc.
in Maynard, Massachusetts

Faint, illegible text at the bottom of the page, possibly bleed-through or a secondary title.



aining, -dains. To ap-
reordain. —pre-or-di-

a prep course. —
2. *British Slang*. The
v. **prepped**, **prepping**.
I attend a preparatory
ion for something.
e. 2. preposition.

g. —ages. To wrap or
ng.

prep. 1. The act or
g made ready before-
inary measures that
rations for the wedding
icine, prepared for a
icipation of a disso-
a consonant tone in
d.

7. Serving or tending,
ich prepares for some-

person who prepares,
or display.
par'-) *adj.* **Abbr.** **prepa-**
inary, introductory.
ration, especially for
In preparing for. See

is, is restricted to situ-
s not of the mere equiva-
s so used informally,
usually private, pre-
in, for public school.
ool."

es. —*tr.* 1. To make
or for some event, oc-
r make by combining
ature; compound
ared for service in the
n (a dissonance or its
To put things or one-
glish *prepara*, from
re, to prepare in ad-
re (see *per-* in Ap-
er *n.*
ite of being prepared;

s. To pay or pay for

l or arranged in ad-
e phrase *malice pre-*
epped, from Middle
om Old French *pour-*
s, from Latin *pr-* +
native of *pendere*, to
ense'*v. adv.*

iso **pre-pon-der-ant-**
ower, importance, or

iving superior powers
nt. See *Synonyms at*

-ated; -ating, -ates;
2. To be greater, in-
e like; predominat-
ns, the latter *seemed*
haic. To be weighed
ponderare '*v. praec.* in
1, from *pondus* (stem
l). —**pre-pon-der-ant-**

rep. *Grammar*. 1. In-
ie relation of a sub-
er substantive. Some
m, and with. 2. Any
uch as in *regard to* or
n, from Latin *praec-*
om *praepone* (past
' *praec.* in front +

in ending a sentence
e awkwardness by
arrogant manner he
eak ending: *He wel-*
te he hoped to spend
al position is the only
much to be thankful
awkward or stilted
eeding, the natural

elating to, composed
ion-al-ly *adv.*
consisting of a prep-
ing adjectival or ad-
wool and written in
ol (adjectival value)

ar. Put before, pre-

fixed. —*n. Grammar*. A word or particle put before another
word. [Late Latin *praepositivus*, from *praepōnere* (past participle
praepositus), to place in front. See *preposition*.] —**pre-**
pos-itive-ly adv.

pre-pos-sess (prē-pō-zēs') *tr. v.* **-sessed, -singing, -singing.** 1. To
preoccupy the mind of to the exclusion of other thoughts or
feelings. 2. To influence beforehand against or in favor of
someone or something; prejudice; bias. 3. To impress favorably
in advance.

pre-pos-sess-ingly (prē-pō-zēs'ing) *adj.* 1. Impressing favorably;
pleasing. 2. *Archaic*. Causing prejudice. —**pre-pos-sess-ingly**
adv. —**pre-pos-sess-ing-ness n.**

pre-pos-sition (prē-pō-zēsh'ōn) *n.* 1. A preconception or
prejudice. 2. The state of being preoccupied with thoughts,
opinions, or feelings.

pre-pos-ter-ous (prī-pōs'tor-əs) *adj.* Contrary to nature, reason,
or common sense; absurd. See *Synonyms at* **foolish**.
[Latin *praeposterus*, "inverted," perverted, absurd: *prae-*,
before + *posterus*, coming after, following, next, from *post*,
after (see *apo-* in Appendix*).] —**pre-pos-ter-ous-ly adv.**
—**pre-pos-ter-ous-ness n.**

pre-po-ten-cy (prī-pō'ten-sē) *n.* The state or condition of being
prepotent; predominance.

pre-po-tent (prī-pō'tent) *adj.* Also **pre-po-tent-tial** (prē-pō-tēn'-
shəl). Greater in power, influence, or force; predominant.
[Middle English, from Latin *praepotēns*, present participle of
praeposuisse, to be very powerful: *prae-* (intensifier) + *posse*, to be
able or powerful (see *pot-* in Appendix*).] —**pre-po-tent-ly adv.**
prep school. *Informal*. A preparatory school (see).

pre-puce (prē-pyūs) *n.* 1. The loose fold of skin that covers
the glans of the penis. Also called "foreskin." 2. A similar
structure covering the glans of the clitoris. [Middle English,
from Old French, from Latin *praepūtium*. See *pu-* in Appen-
dix.*] —**pre-pu-tial** (prī-pyūs'eshəl) *adj.*

pre-Raph-a-el-ite (prē-rā-fē-lit', prē-rā-fē-) *n.* A painter or
writer belonging to or influenced by the pre-Raphaelite Broth-
erhood, a society founded in 1848 by Rossetti and others to
advance the style and spirit of Italian painting before Raphael.
—*adj.* Of, pertaining to, or characteristic of the pre-Raphael-
ites. —**pre-Raph-a-el-ite-ism n.**

pre-req-u-isite (prē-rēk'wō-zit) *adj.* Required as a prior condi-
tion to something. See *Synonyms at* **necessary**. —*n.* That
which is prerequisite.

Usage. The appropriate prepositions after this word are *to*
(for the adjective) and *of* (for the noun).

pre-rog-a-tive (prī-rōg'ō-tiv) *n.* 1. An exclusive right or privi-
lege held by a person or group, especially a hereditary or official
right. 2. Any characteristically exclusive right or privilege.
3. A natural gift or advantage making one superior. 4. *Obso-*
lete. Priority or pre-eminence; superiority. —*See Synonyms at*
right. —*adj.* Of, arising from, or exercising a prerogative.
[Middle English, from Old French, from Latin *praerogātiva*
(*centuria*), "(century) chosen to vote first," from *praerogātus*,
asked to vote first, from *praerogāre*, to ask before others: *prae-*,
before + *rogāre*, to ask (see *reg-* in Appendix*.)]

pres. 1. present (time). 2. president.

Pres. President.

pres-age (prēs'ij) *n.* 1. An indication or warning of a future
occurrence; omen; portent. 2. A feeling or intuition of what is
going to occur; presentiment; foreboding. 3. Prophetic signifi-
cance or meaning. 4. *Rare*. A prediction. —*v.* **pres-age** (prī-
sāj') **-aged, -aging, -ages, -tr.** 1. To indicate or warn of in
advance; portend. 2. To have a presentiment of. 3. To foretell
or predict. —*intr.* To make or utter a prediction. —*See Syn-*
onyms at **foretell**. [Middle English, from Latin *praesagium*,
foreboding, from *praesagere*, to perceive beforehand: *prae-*,
before + *sagere*, to perceive (see *sag-* in Appendix*.)] —**pres-**
age-ful (prī-sāj'fəl) *adj.*

pres-by-ol-ly (prēz'bē-ō-lē, prēs'-) *n.* The inability of the eye
to focus sharply on nearby objects, resulting from hardening of
the crystalline lens with advancing age. [New Latin: Greek
presbys, old man (see *per-* in Appendix*) + *-opia*.] —**pres-**
by-opic (ōp'ik) *adj.*

pres-by-ter (prēz'bō-tor, prēs'-) *n.* Ecclesiastical. 1. In the early
Christian church, an elder of the congregation. 2. In various
hierarchical churches, a priest: 3. *Presbyterian Church*. a. A
teaching elder. b. A ruling elder. [Late Latin, an elder, from
Greek *presbuteros*, a priest, "older," comparative of *presbus*,
old man. See *per-* in Appendix*.]

pres-by-ter-ate (prēz-bī't-or-īt, -ō-rāt', prēs-) *n.* 1. The office of
a presbyter. 2. The body or order of presbyters.

pres-by-ter-i-al (prēz'bō-tēr'ē-əl, prēs'-) *adj.* Of or pertaining to
a presbyter or the presbytery. —**pres-by-ter-i-al-ly adv.**

pres-by-ter-i-an (prēz'bō-tēr'ē-ən, prēs'-) *adj.* 1. Of or pertain-
ing to ecclesiastical government by presbyters. 2. *Capital P*. Of
or pertaining to a Presbyterian Church. —*n.* *Capital P*. A
member or adherent of a Presbyterian Church. —**pres-by-ter-**
ri-an-ism n.

Presbyterian Church. Any of various Protestant churches
governed by presbyters and traditionally Calvinist in doctrine.
pres-by-ter-y (prēz'bō-tēr'ē, prēs'-) *n.*, *pl.* -ies. 1. *Presbyterian*
Church. a. A court composed of the ministers and representa-
tive elders of a particular locality. b. The district represented by
this court. 2. Presbyters collectively. 3. Government of a
church by presbyters. 4. The section of the church reserved for
the clergy. 5. *Roman Catholic Church*. The residence of a
priest. [Middle English *presbytery*, from Late Latin *presby-*
terium, a council of presbyters, from Greek *presbuterion*, from
presbuteros, priest, **PREBYTER**.]

pre-school (prē'skōl') *adj.* Of or pertaining to a child of
nursery-school age. —**pre-school'er n.**

pre-sci-ence (prē'shē-əns, prēsh'ē-) *n.* Knowledge of actions or
events before they occur; foreknowledge; foresight.

pre-sci-ent (prē'shē-ənt, prēsh'ē-) *adj.* 1. Of or pertaining to
prescience. 2. Possessing prescience. [Latin *praesciens*, pres-
ent participle of *praescire*, to know beforehand: *prae-*, before +
scire, to know (see *skol-* in Appendix*.)] —**pre-sci-ent-ly adv.**

pre-scind (prī-slīnd') *v.* **-scinded, -scinding, -scinding.** —*tr.* To
separate or divide in thought; consider individually. Used with
from. —*intr.* To withdraw one's attention. Used with *from*.
[Latin *praescindere*, to cut off in front: *prae-* in front +
scindere, to cut off (see *skol-* in Appendix*.)]

Pres-cott (prēs'kət), **William**. 1726-1795. Commander of the
Continentalists at Bunker Hill. Grandfather of William Hickling
Prescott.

Pres-cott (prēs'kət), **William Hickling**. 1796-1859. American
historian of Spain and the Spanish conquests in the Americas.
Grandson of William Prescott.

pre-scribe (prī-skrīb') *v.* **-scribed, -scribing, -scribes.** —*tr.* 1. To
set down as a rule or guide; ordain; enjoin. 2. *Medicine*. To
order or recommend the use of (a drug or other therapy).
—*intr.* 1. To establish rules, laws, or directions. 2. *Medicine*.
To order or recommend a remedy on the grounds of prescrip-
tion. b. To become invalidated or unenforceable by the process
of prescription. [Middle English *prescribere*, to claim by such
right, from Latin *praescribere*, to write at the beginning, *pre-*
scribe: *prae-*, before, in front + *scribere*, to write (see *skol-*
in Appendix*).] —**pre-scrib'er n.**

pre-script (prē'skrīpt') *n.* Something prescribed, especially a
rule or regulation of conduct. —*adj.* (prē'skrīpt, prī-skrīpt').
Established as a rule; set down; prescribed. [Latin *praescriptum*,
from *praescribere* (past participle *praescriptus*), **PRESCRIBE**.]

pre-scrip-ti-ble (prī-skrīpt'ē-ə-bəl) *adj.* Capable of, requiring, or
derived from prescription. —**pre-scrip-ti-ble-ly n.**

pre-scrip-tion (prī-skrīpt'shən) *n.* 1. The act of prescribing.
b. That which is prescribed. 2. *Medicine*. a. A written instruc-
tion by a physician for the preparation and administration of a
medicine. b. A prescribed medicine. c. An ophthalmologist's
or optometrist's written instruction for the grinding of correc-
tive lenses. 3. A formula directing the preparation of anything.
4. *Law*. a. The process of acquiring title to property by reason
of uninterrupted possession of specified duration. Also called
"positive prescription." b. The limitation of time beyond which
an action, debt, or crime is no longer valid or enforceable. Also
called "negative prescription." [Middle English *prescription*,
from Old French *prescription*, from Latin *praescriptio*, a writing
in front, from *praescribere*, **PRESCRIBE**.]

pre-scrip-tive (prī-skrīpt'iv) *adj.* 1. Sanctioned or authorized
by long-standing custom or usage. 2. Making or giving injunc-
tions, directions, laws, or rules. 3. *Law*. Acquired by or based
upon uninterrupted possession. —**pre-scrip-tive-ly adv.**

pre-sell (prēs'sēl') *tr. v.* **-sold (-söld'), -selling, -sells.** To promote
(a product not yet on the market).

pres-ence (prēs'əns) *n.* 1. The state or fact of being present.
2. Immediate proximity in time or space. 3. a. The area im-
mediately surrounding a great personage, especially a sovereign
granting audience. b. A person who is present. 4. A person's
manner of carrying himself; bearing. 5. A supernatural influ-
ence felt to be nearby. —*See Synonyms at* **bearing**.

presence of mind. Ability to think and act efficiently, espe-
cially in an emergency.

pres-ent (prēz'ənt) *n.* 1. A moment or period in time percep-
tible as intermediate between past and future; now. 2. *Abbr. pr.*
pres. Grammar. a. The present tense. b. A verb form in the
present tense. 3. *Plural. Law*. The document or instrument in
question: *be it known by these presents*. —*adj.* 1. Being, per-
taining to, or occurring at a moment or period in time con-
sidered as the present. 2. Being at hand. 3. *Obsolete*. Alert to
circumstances; attentive. 4. *Obsolete*. Readily available; im-
mediate. 5. *Abbr. pr., pres. Grammar*. Denoting a verb tense or
form that expresses current time. [Middle English, from Old
French, from Latin *praesens*, present participle of *praeesse*, to
be before one, be present: *prae-* in front of + *esse*, to be (see
es- in Appendix*).]

pre-sent (prī-zēnt') *tr. v.* **-sented, -senting, -sents.** 1. a. To in-
troduce, especially with formal ceremony: *This Brazilian*
king... was brought up to London and presented to King Henry
VIII (Richard Hakluyt). b. To introduce (a girl) to society
with conventional ceremony: *she also was presented that year*
at the International Debutante Ball (New York Times). 2. To
bring before the public: *present a play*. 3. a. To make a gift or
award of: *present a medal*. b. To make a gift to; bestow for-
mally: *present the college an endowment*. 4. To offer for consid-
eration; display: *present one's credentials*. 5. To offer for consid-
eration. 6. To salute with or aim (a weapon). 7. *Ecclesiastical*. To rec-
ommend (a clergyman) for a benefice. 8. *Law*. a. To offer to a
legislature or court for consideration. b. To bring a charge or
indictment against. —*See Synonyms at* **offer**. —*n.* **pre-sent**
(prēz'ənt). Something presented; a gift. [Middle English
praesentem, from Old French *praesentem*, from Latin *praesentare*,
from *praesens*, **PRESENT** (adjective).] —**pre-sent'er n.**

pre-sent-a-ble (prī-zēnt'ə-bəl) *adj.* 1. Capable of being given,
displayed, or offered. 2. Fit for introduction to others. —**pre-**
sent-a-ble-ly adv.

pre-sent-a-ble-ness n. —**pre-sent-a-ble-ly adv.**

pre-sen-ta-tion (prēz'ən-tā'shən, prēz'ən-) *n.* 1. *n.* The act of
presenting or offering for acceptance or approval. b. The state



pre-Raphaelite
Detail of a painting by
Dante Gabriel Rossetti

t tight/th thin, path/th this, bathe/ü cut/ür urge/v valve/w with/y yes/z zebra, size/zh vision/ö about, item, edible, gallop, circus/
ä Fr. ämi/äc Fr. feu, Ger. schön/ü Fr. tu, Ger. über/kh Ger. ich, Scot. loch/N Fr. bon. *Follows main vocabulary. †Of obscure origin.

/j judge/k kick/l lid/
sauce/sh ship, dish/

EXHIBIT 13

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

TELCORDIA TECHNOLOGIES, INC.,)	
)	
Plaintiff,)	
)	
v.)	Civil Action No. 04-874 GMS
)	
ALCATEL USA, INC.,)	
)	
Defendants.)	
<hr/>		
TELCORDIA TECHNOLOGIES, INC.,)	
)	
Plaintiff,)	
)	
v.)	Civil Action No. 04-875 GMS
)	
LUCENT TECHNOLOGIES, INC.,)	
)	
Defendants.)	
<hr/>		
TELCORDIA TECHNOLOGIES, INC.,)	
)	
Plaintiff,)	
)	
v.)	Civil Action No. 04-876 GMS
)	
CISCO SYSTEMS, INC.,)	
)	
Defendants.)	
<hr/>		

ORDER

WHEREAS, on July 16, 2004, the plaintiff, Telcordia Technologies, Inc. (“Telcordia”), filed the above-captioned patent infringement actions against Alcatel USA, Inc. (“Alcatel”), Lucent Technologies, Inc. (“Lucent”), and Cisco Systems, Inc. (“Cisco”) (collectively, the “defendants”);

WHEREAS, on February 17, 2006, the parties submitted a Final Joint Claim Chart (the “Chart”) (D.I. 98);

WHEREAS, upon inspection of the Chart, the court has discovered that many of the defendants’ proposed constructions for U.S. Patent Nos. Re. 36,633 and 4,835,763¹ are not constructions but, rather, arguments that the claim limitations at issue are indefinite for failure to satisfy the requirements of 35 U.S.C. § 112(2); and

WHEREAS, the court does not permit summary judgment arguments, including indefiniteness arguments, during the claim construction phase of the litigation;

IT IS HEREBY ORDERED that:

1. The court will not entertain indefiniteness arguments during the *Markman* Claim Construction hearing.
2. The defendants shall prepare their arguments consistent with this Order.

Dated: April 21, 2006

/s/ Gregory M. Sleet
UNITED STATES DISTRICT JUDGE

¹ Telcordia’s action against Alcatel does not include U.S. Patent No. 4,835,763.

EXHIBIT 14

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

NETRATINGS, INC.,)
)
 Plaintiff,)
)
 v.) Civil Action No. 05-314 GMS
)
 COREMETRICS, INC.,)
)
 Defendant.)

ORDER

WHEREAS, on May 19, 2005, the plaintiff, NetRatings, Inc. filed the above-captioned patent infringement action against Coremetrics, Inc.;

WHEREAS, on April 3, 2006, the parties submitted a Joint Claim Chart (the “Chart”) (D.I. 47);

WHEREAS, upon inspection of the Chart, the court has discovered that many of the defendant’s proposed constructions for U.S. Patent No. 6,108,637 are not constructions but, instead, arguments that the claim limitations at issue are indefinite, nonenabled, and/or invalid for failure to disclose the best mode; and

WHEREAS, the court will not permit summary judgment arguments, including indefiniteness, enablement, or invalidity arguments, during the claim construction phase of the litigation but, rather, rely on the defendant’s claim construction briefs for those arguments;

IT IS HEREBY ORDERED that:

1. The court will not entertain indefiniteness, enablement, or invalidity arguments during the *Markman* Claim Construction hearing.
2. The defendants shall prepare their arguments consistent with this Order.

Dated: June 7, 2006

/s/ Gregory M. Sleet
UNITED STATES DISTRICT JUDGE