

Recommendation Systems:
Decision Support for the Information Economy

Edward A. Stohr

Sivakumar Viswanathan

May 1998

Working Paper Series
Stern #IS-98-17

Recommendation Systems: Decision Support for the Information Economy

Edward A. Stohr
and
Sivakumar Viswanathan

Information Systems Department
Leonard N. Stern School of Business
New York University
44 West 4th Street
New York, NY 10012

To be published in Kenneth E. Kendall (ed.), "Emerging Information Technologies.", Thousand
Oaks, CA: Sage Publication, Inc., 1999.

INTRODUCTION

It is five years hence, the dawn of the new millenium. The World Wide Web serves two hundred million people and is inhabited by perhaps a billion software agents. A large part of the world economy is on-line. Organizations are networked in a constantly changing kaleidoscope of relationships - some fleeting others mutating slowly. Competition is for the intelligent and agile - those organizations that have mastered the arts of discovering new value propositions, nurturing customer relationships, and executing their core functions rapidly and faultlessly. Human beings in the developed nations are totally dependent on the electronic world for their education, entertainment, communication and commerce.

Is such a world possible? Simple extrapolation of current exponential growth trends indicates that the above human and software agent population figures are not unreasonable. The feasibility of such a vision, however, rests on issues of system stability and human benefits. Can an electronic economy making millions of decisions per minute and executing billions of transactions and transferring trillions of dollars per day remain stable? Or will there be an electronic world financial crisis that will make the recent Asian crisis seem like a minor blip? Will the evolving electronic society actually benefit mankind? Will it create an elite of whiz-kid millionaires controlling vast electronic financial resources and a huge underclass of the less fortunate or less able? We are already experiencing an almost unbearable flood of information and communication requirements that sap our time, make leisure more difficult, and our time more and more valuable. How can we realize the benefits of the increased productivity that we feel instinctively can be provided by the information revolution?

The questions raised in the last paragraph, can probably only be answered in practice as new markets, organizational forms and life styles evolve from the maelstrom of electronic commerce. In this chapter, we will concentrate on only one necessary condition for the existence of such an electronic world. Our thesis is that we need new technologies to help individuals and firms cope with the evolving information ecology - coping technologies to manage the information producing technologies (which if unchecked could actually become disabling technologies). The emphasis in information technology research and practice must turn from producing information to controlling and managing it. Shaping the information into forms that are more useful and accessible, condensing it, filtering out noisy elements, finding useful nuggets of information on demand, analyzing them, and creating new knowledge - using every means at our

disposal to make the new information universe friendly, understandable, and useful. The technologies for doing this are exciting and challenging. Some have been around since the early days of computer and management science, others are emerging in response to the new electronic realities.

The technologies fall into four groups (see Figure 1). The technologies for discovering information: search engines, directories, electronic markets and electronic auctions. The technologies for controlling and restricting the flow of information to which we are subject: filtering and alerting systems. The technologies for understanding information: knowledge representation, visualization, data mining, and the tools of statistics and management science. The technologies to assist decision making: recommender systems and "electronic butler" systems [23]. The first three categories of technology - those for discovering, controlling and understanding information - provide support for the fourth category - that of decision support and decision making.

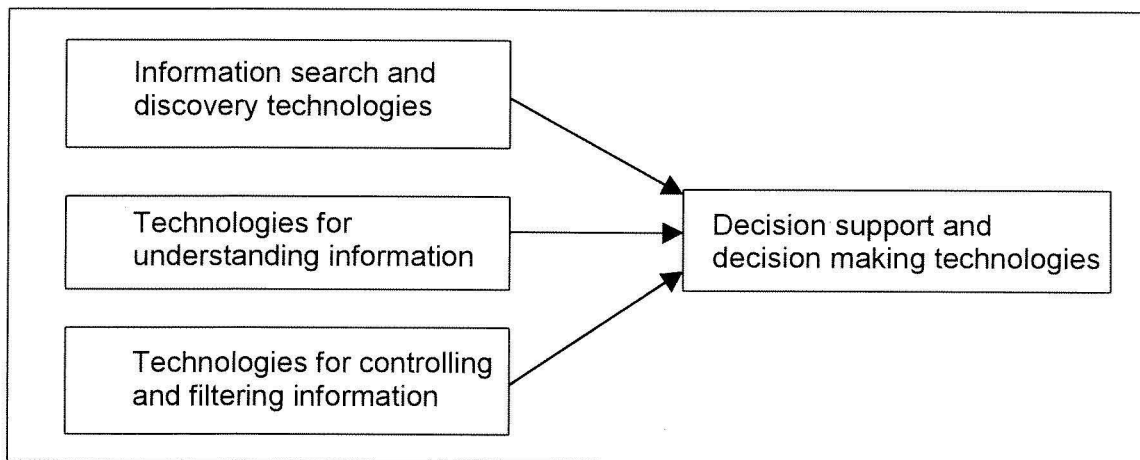


Figure 1

A Classification of Technologies for Coping with the Information Explosion

In this chapter, we will concentrate on the fourth set of technologies and within that only on a new class of systems that have been called recommender systems. Our objective is to provide an overview of recommender systems and their role in the information economy. We look at where they fit in the panoply of technologies that have been developed to support decision making in electronic commerce, their technical foundations, managerial implications, and issues that surround their deployment and adoption.

RECOMMENDATION SYSTEMS AS A DECISION TECHNOLOGY

Recommendation systems provide information about the relative merits of alternative courses of action. In everyday life, when faced with a choice for which we lack adequate information, we often turn to various guides such as Consumer Choice magazine, Zagat's restaurant guide, knowledgeable friends, experts, and so on. The function of these guides is to increase the probability that we will be satisfied with the results of our decision making. Recommendation systems are electronic versions of such everyday systems.

There is a subtle definitional issue. As originally conceived, and by common usage, recommender systems are automated systems in which people (recommenders) provide recommendations as inputs, which the systems then aggregate and direct to appropriate recipients [16]. This is narrower than the definition of the previous paragraph, which essentially says that recommender systems make recommendations (by any means.) We prefer the broader definition: first, because it is more goal-oriented and second, because the narrow definition seems unnecessarily restrictive. Strictly interpreted, for example, the narrow definition would not include the use of intelligent software agents as recommenders. Nor would it include information filtering systems that make recommendations based on content analysis in which a user's interest profile is matched with document content profiles. Worse still, many recommender systems are really hybrids in which the source of the recommendations is both human and non-human (collaborative filtering plus content analysis, for example). To distinguish between the broad and narrow definitions of recommender systems, we will call the former "recommendation" systems. Because this is a new area for computer support, it is worth while taking some pains to make the definitions as clear as possible. We move from the general to the particular (see figure 2.).

The objective of *decision technologies* in general is to overcome the limits of our bounded rationality - to help us make better decisions, to make them more quickly, and with less expenditure of effort. *Decision making technologies* automate the decision making process and remove the human element. Automated inventory systems are a prime example. More recently, [23] has suggested that automated decision making be extended to include everyday activities such as shopping. He has proposed "electronic butler" services that use the past history of user purchases to infer a subset of purchases that can be made automatically without user consultation.

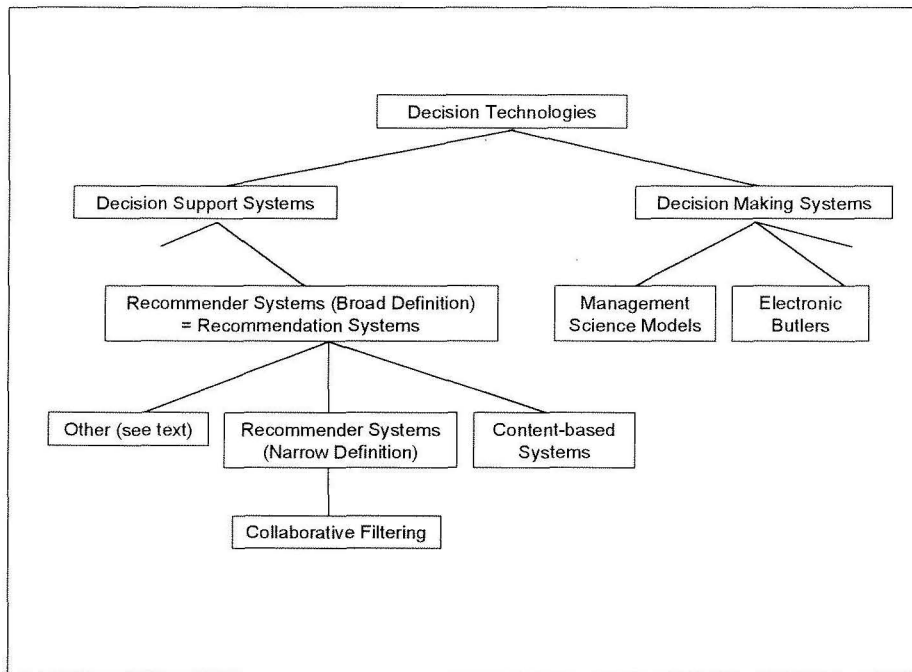


Figure 2
Partial Classification of Decision Technologies

Decision Support Systems (DSS) were originally conceived as an alternative to computerized decision making systems for use in situations in which the human inputs were necessary because the decision situation was not highly structured (programmable) and therefore, required human judgement and intuition [5]. The role of the computer was to provide information via databases or models to help in the decision making process [1]. An interesting consequence of the new information economy is that a decision situation can become unstructured as a result of a glut of information and the speed at which decisions have to be made. Users then need computer support, not only because of the complexity of a given decision (the original DSS concept), but also, because they lack the processing speed and power to cope with an overload of information and a myriad of decisions that demand attention.

Recommendation systems, systems that evaluate choices, fall in this category of decision support. Typical uses of recommendation systems have been to suggest the information that a decision maker might need (filtering systems), and to rank films,

restaurants, books, and so on¹. These are small decisions in a way, but demanding in the aggregate. However, we do not mean to restrict recommendation systems to low or medium matters of importance to individuals or organizations. Systems to assist users in decisions of major consequence, such as equipment purchases for corporations and home or automobile purchases for individuals, can also be envisaged as outlined below. Recommendation systems may use all or any information from whatever source and make use of very different processing algorithms to form their recommendations.

Recommender systems (narrow definition) on the other hand, perform the same function as recommendation systems, but involve other human beings as recommenders. The recommenders may or may not explicitly collaborate with the recipients of the information because the recommenders and recipients may not be known to each other [16].

Finally, *collaborative filtering systems* are special purpose recommender systems that filter or restrict the information delivered to recipients. They are narrower in scope than recommender systems because the latter can suggest new items of interest rather than just filter out those that are presumed to have little or no interest [16]. Typically, collaborative filtering systems have been used to filter information from Internet Usenet groups (as in Grouplens [9]), provide recommendations on books (as in Amazon.com), etc. Collaborative filtering techniques can differ in their implementation and approach, but all incorporate the use of individuals to annotate or recommend items to others. Of course, the same individual can alternate between the roles of recommender and recipient at different times.

RECOMMENDATION SYSTEMS WITHIN A FRAMEWORK FOR ELECTRONIC COMMERCE

The World Wide Web has dramatically altered the availability of information for users, but most of this information is useless or of poor quality. Given the magnitude and diversity of content and its uncertain quality, familiar search engines, such as Lycos, Yahoo and Alta Vista, have proved inadequate. Typically, these search engines employ a fairly traditional information retrieval (IR) paradigm augmented by software agents (spiders) that continuously examine and index the content of millions of web sites. The metamorphosis of the World Wide Web from a *digital library* (where the focus was on

¹ Many existing examples of recommendation systems fit the narrow definition of recommender systems.

retrieval) into an *electronic marketplace* (where the focus now is on transactions) alters the rules of the game, calling for a new perspective and a broad range of new capabilities, institutions and control mechanisms. The following framework provides some pointers for future directions and some insights into the role of recommendation systems in the information economy.

The framework is based on the concept of *intermediation* in an economy. Just as the traditional economy is populated by physical intermediaries in various forms, who serve to facilitate transactions, the information economy requires the service of ‘electronic intermediaries’ serving similar roles² on the Web. Intermediaries, by nature of their role-specialization, help to reduce the risks and uncertainties that plague transactions in a traditional economy. In addition, intermediaries add value by engaging in activities that help reduce coordination costs and provide economies of scale and scope. The information economy could, similarly, be greatly benefited by intermediaries who can help facilitate transactions.

The framework in Figure 3 is an adaptation and expansion of that contained in [7]. The framework proposes a hierarchy of ten levels of functionality. While some of the levels of functionality primarily focus on the needs of customers, others serve the needs of suppliers/ manufacturers. Note that these levels are not necessarily conceptualized as strict levels in a hierarchy, i.e., we do not claim that higher levels of functionality, by themselves, add more value than those below them. While the higher levels of functionality indicate potentially greater benefit, they also rank higher in terms of complexity of implementation and execution besides requiring human intervention and consensus among various players. The higher levels of functionality, being inherently more complex in nature, have been slower to evolve. Thus the various levels in the framework may be conceptualized as stages in the evolution of e-commerce systems. Web-based agents typically focus on one or more levels of functionality and strive to gain efficiencies in these. A more complete description of the framework and its implications for the design of web-based agents and for electronic commerce in general is contained in [21].

² Since we are primarily concerned with ‘electronic intermediaries’ in an *information economy*, we restrict our discussion to intermediary functions that lend themselves to automation.

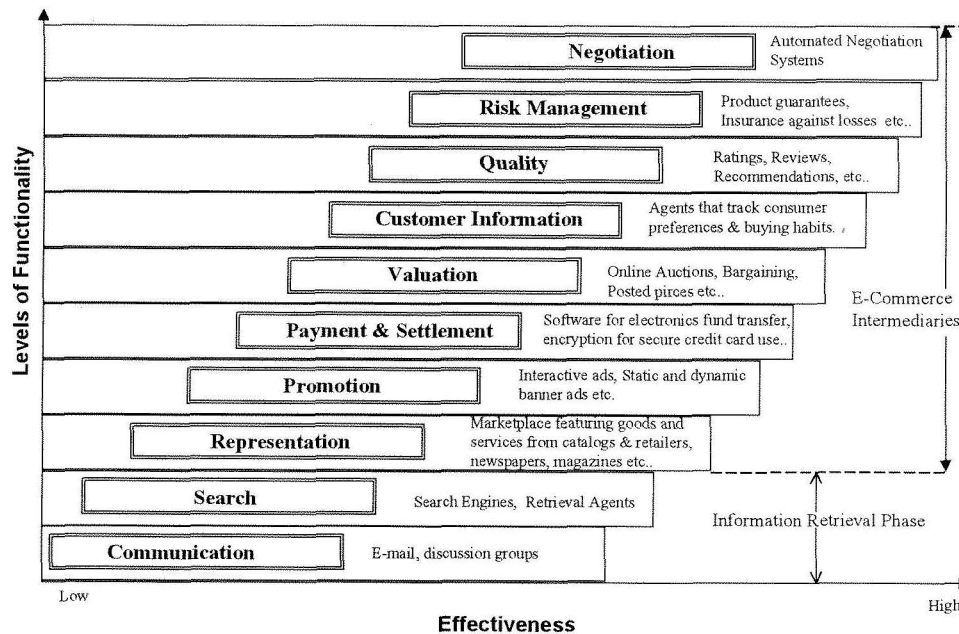


Figure 3
A Value Framework for Web-based Agents

The above framework highlights the role of recommendation systems in electronic commerce. Recommendation systems are software agents that attempt to incorporate the three functions of search and retrieval, eliciting customer information, and signaling quality. In addition, because they convey quality information, they help mitigate the risk of executing transactions on the Web.

The combination of the three functions of search and retrieval, eliciting customer information, and signaling quality in recommendation systems makes them ideally suitable for "one-to-one marketing" applications. In one-to-one marketing, firms seek to learn and satisfy the unique needs of each individual customer [15]. For example, Individual Inc.'s First! Service allows it to compete with wire, clipping and information retrieval services [<http://www.individual.com/>]. Clients such as MCI Telecommunications, McKinsey & Co., and Avon products, use the service to have customized information delivered to executives via fax, e-mail, Lotus Notes, or their corporate Intranets. The executives first provide simple descriptions of the kinds of items in which they are interested. The system then fine tunes these "user profiles" by asking the executives to rate each received article as being of high, medium or low relevance. Over time, the ratio of medium plus highly relevant articles to the total number of articles received by

the executives, is reported to be as high as 80% [15]. Firefly (<http://www.firefly.com/>) is another example of a commercial recommendation system that facilitates one-to-one marketing. The Firefly software enables businesses to create user-centric applications and services that manage personal profiles and personalize services for individual users. For example, Filmfinder (<http://www.filmfinder.com/>), a site that recommends movies, is enabled by Firefly. Other companies that have formed partnerships with Firefly include Barnes and Noble, Virtual Emporium, and Yahoo.

Other applications of recommendation systems on the Web include: personalized recommendations of URLs, filtering of Usenet articles, blocking access to undesirable Web sites, one-to-one marketing of banking services, and shopping services for music, video and books [16]. Some of these applications are discussed in more detail below. Within an organization, recommendation systems will be increasingly used in organizational learning applications in which employees share knowledge by commenting on and rating various company products, sales leads and practices [20].

INFORMATION RETRIEVAL – THE UNDERLYING TECHNOLOGY FOR RECOMMENDATION SYSTEMS

To understand recommendation systems, it is useful to compare them with the information retrieval systems from which they have evolved. In this section, we provide a broad overview of IR as a background to the framework for recommendation systems that we develop in the next section.

An information retrieval system helps users find documents that best satisfy their need for information, or helps them obtain information from knowledge sources for the purposes of problem management [3]. The term ‘information retrieval’ is primarily associated with text retrieval. The three major areas of concern in IR Research are depicted in Figure 4. These are: representation of the individual’s information needs; representation of the meaning of the texts/documents, and comparison of the two to find the most suitable match.

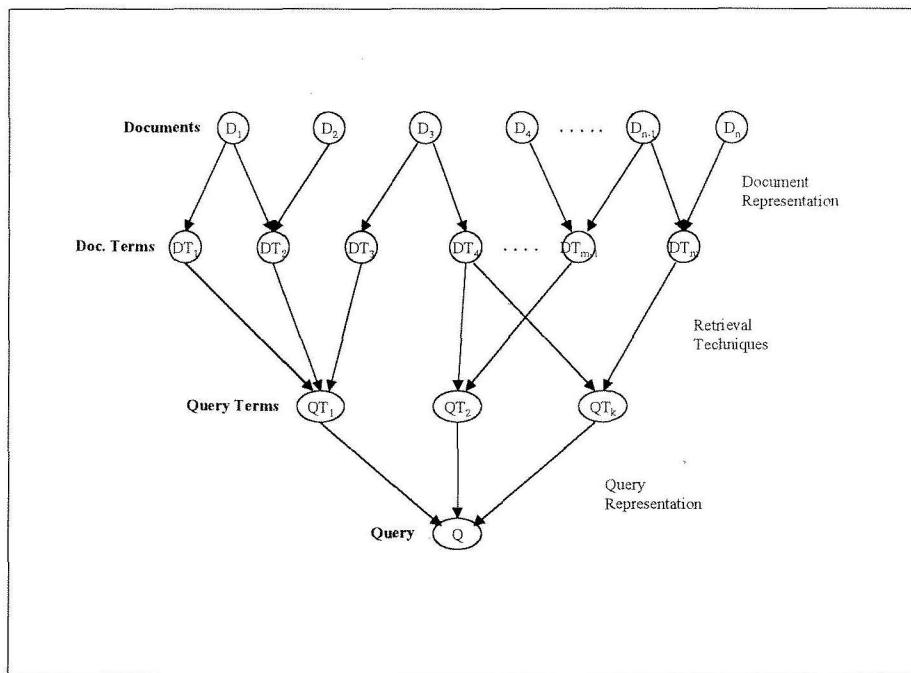


Figure 4
Basic Inference Network

- Information Needs Representation – The information needs of a user are usually represented by means of a query consisting of a simple term or a set of terms. The query terms are usually keywords or phrases involving Boolean representation. However, given the richness of natural language, query terms are usually not accurately indicative of the user’s true meaning.
- Document Representation – Documents or texts produced by different sources are indexed (manually or automatically) to extract terms that are best representative of the documents. As in the case of queries, these document terms too are not accurately indicative of the actual document meaning. Various indexing techniques (for example Probabilistic Indexing and TF-IDF [18]) have been developed in an attempt to more accurately characterize the information content of the document.
- Retrieval Techniques – The third area of concern is the comparison of the query terms with the document terms, based on which the relevant documents are retrieved. Most of the major IR models have been primarily concerned with this process of comparison.

It is inherently difficult to satisfy specific user needs with a document database

indexed by librarians to serve the needs of a general audience. "Associative" search techniques based on probability, statistics, set theory and logic (for example, Boolean models, Vector Space Models and Probabilistic Retrieval Models [17]), enable the retrieval of documents that are "close" in some sense to the user's query. These same measures of closeness can narrow the user's search for relevant documents by providing a system-determined *relevance rating* or *ranking* of retrieved articles in terms of their "relevance" to the user's request. In addition, "relevance feedback" techniques - continuously improving queries by asking the user to rate the relevance of retrieved articles - can greatly improve the performance of IR systems. IR systems that provide relevance ratings are a specialized form of recommendation system as we have defined the term.

Current IR systems depend on the 'best-match' principle, i.e., given a query, the best possible system response is the text whose representation most closely matches it [4]. The 'best-match' principle in turn depends on the assumption of equivalence between the *expression of need* and the document text in that it treats the representation of the need as a representation of the document that is ideal for resolving that need. The best-match principle looks first for a document, which is just like the expression of need; that is, which is functionally equivalent to it [4]. If we know precisely what the user wants and if we know what documents best satisfy this need, then the problem reduces to a simple matching process. Unfortunately, "noise" is inherent in each of the IR components in Figure 4. Users generally cannot specify precisely what information is needed to resolve a particular problem, documents are represented imprecisely by the document terms, and, while many different algorithms have been proposed, the retrieval process may be unsatisfactory.

Precision and *recall* are the two most common measures for evaluating IR effectiveness. *Precision* is defined as the percentage of items retrieved in a search that are relevant to the query, while *recall* is defined as the percentage of relevant items that are retrieved in a search. While one would like the IR system to consistently score high on both measures, there is an inherent trade-off between the two. The concept of *relevance* has in itself, been a source of problems. Although, there has been a lack of consensus regarding the concept of relevance, a "system's view" of relevance has dominated most of IR research. As stated by Saracevic (1976) [19], "*the system's view of relevance was a result of thinking that relevance is mostly affected by the internal*

aspects and manipulations of the system. Relevance was conceived in terms of indexing, coding, classification, linguistic manipulations, file organization, and eventually question analysis and searching strategies. This thinking led to development of a myriad of schemes, and to attention to input processing and manipulation almost to the exclusion of other aspects". Saracevic emphasized the need for a view of relevance that included the concept of pragmatic usefulness or "pertinence" to the needs of the user.

The problems and limitations of IR systems become acute as we move to the domain of Web-based commerce. Most of the early attempts at information retrieval on the Web (for example search engines and directories) have built upon the models of IR, notwithstanding the fact that these systems performed best in narrow domains where the information was relatively well structured and homogenous. The Web, though, is a very different information space, with vast differences in the structure, quantity and quality of content. In particular, the development of electronic commerce, where the focus is on transactions involving consumer durables as well as informational items, demands a much higher level of functionality and a quite different view of "relevance."

RECOMMENDATION SYSTEMS ARCHITECTURE

The schematic in Figure 5, provides a framework for recommendation systems illustrating both their basic components and the change in perspective that differentiates them from traditional IR systems. The dotted lines in the figure encompass the areas of concern of IR systems (see Figure 5).

The major differences between IR and recommendation systems are first, that the objects of interest may be general items (goods) as well as informational items, and second, that the pertinence, or relevance of the retrieved items to the user's actual needs (utility), is paramount. These two differences imply capabilities that are not found in IR systems.

Returning to our earlier discussion, recommendation systems provide information about the relative merits of alternative choices or courses of action. The *objects of choice* can be informational, such as articles, books, and web sites. In this case the recommendation system might produce a relevance ranking for the user query as in IR systems, or choose from a large population of information items only those that should be of importance to the user, as in filtering systems. Alternatively, the objects of choice

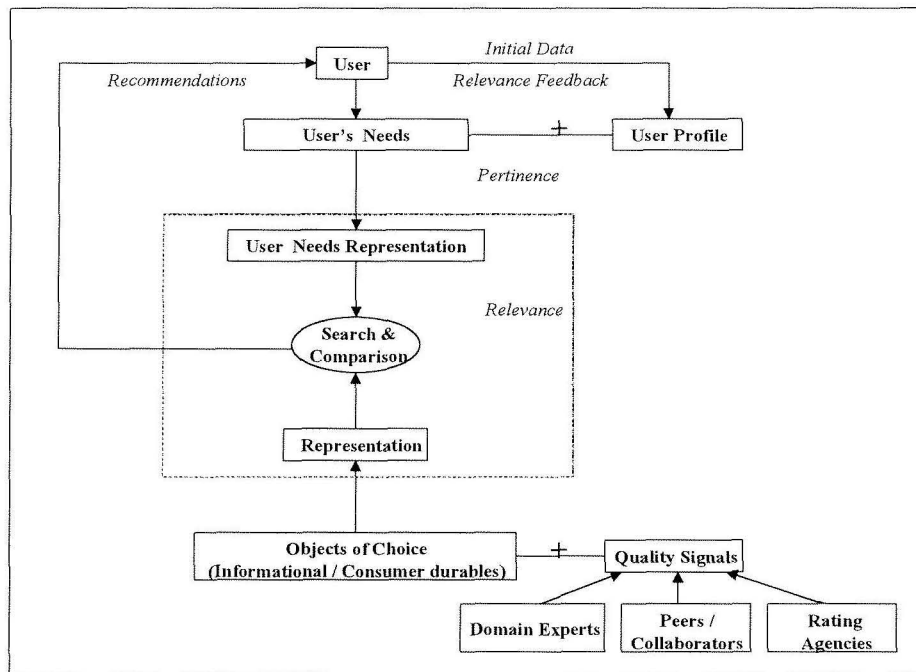


Figure 5
The Architecture for Recommendation Systems

could be durables such as cars and houses; consumable items such as plays, movies and concerts; or abstract entities such as stocks and other investment vehicles. In the sequel, the various alternatives considered by the recommendation system will be said to belong to its "choice domain." The choice domain may be constructed on the fly in answer to a particular user request, or it might be an index or database that is maintained by the system and continuously updated.

Given the broad range of possible applications of recommendation systems, it is of interest to see what minimal set of functions a recommendation system must perform. First, a recommendation system, like an IR system, must have *search capabilities* in order to access and process information about the available choices. In some cases, recommendation systems have been linked to WWW search engines such as Alta Vista [<http://altavista.digital.com/>]; in other cases, they have accessed and processed information in Usenet newsgroups. More generally, the recommendation system might access information in web-based shopping malls, on-line auctions, or the homepages of car manufacturers.

Second, the recommendation system must ensure that the relevant information is available for each alternative in the choice domain. For information items, the relevant information might be annotations and rankings by other users (the recommenders) or a set of document terms as in a traditional IR application; for durable goods, the information on each alternative will usually consist of a list of attribute values including the price of the good, and so on.

Third, the evaluation system must have some notion of user preferences. In most systems, a *user profile* will be generated by asking the user to rank some typical alternatives or to state their preferences and requirements (for example for certain classes of subject matter.) The *user profile* is usually updated on a continuous basis using *relevance feedback techniques*. In other systems, those using rating agencies, it is tacitly assumed that the quality rankings of third party experts can be used as a proxy for the user's own preferences.

Fourth, the recommendation system must be able to *score the alternatives* in its choice domain by assigning them an ordinal rank or cardinal rating that will signal their relative desirability to the user. The scoring mechanisms employed in recommendation systems differ depending on the source of expertise that is used to rank the alternatives as explained in the next section.

Finally, the recommendation system must be able to present its results in an understandable and convenient fashion to the user. Again, there are many alternatives, depending on the type of scoring system used and the application domain. For example, in an IR system, information items are generally ranked in terms of their predicted relevance to the query - and this may be the only choice-oriented information offered by the system. However, especially, in collaborative systems, much more elaborate information may be presented to the user, For example, in Amazon.com the user may read the full text of book critiques by other users.

These five sets of capabilities are necessary in any recommendation system. Namely, the system must be able to seek out possibilities for user choice, generate or maintain relevant information about each alternative in its choice domain, elicit or infer the preference profile of each of its users, evaluate the choices in terms of the user's profile, and present its recommendations to the user. While each recommendation system has these capabilities, there are major differences between them in philosophical approach, which will now be sketched.

CLASSIFICATION AND EXAMPLES OF RECOMMENDATION SYSTEMS

Recommendation systems can be classified according to the source of the knowledge, or expertise, on which the system bases its recommendations. The major alternatives are set out in Table 1 together with some example application domains and references to actual systems that have adopted each approach. As in real life, there are four sources of such expertise. The users themselves (utility approach), information about user preferences and about the choices (content-based systems), the recommendations of a peer group (collaborative systems), and the opinions of third party experts or rating agencies (third party expertise). Many existing Web-based systems combine the content-based and collaborative approaches (hybrid systems). In this section, we discuss each of these approaches, provide examples, and overview their relative advantages and disadvantages.

Recommendation System Technique/Philosophy	Source of System Expertise - Who is the Recommender?	Example Choice Domains	Example System
Utility estimation	User introspection	Big ticket items such as houses and cars	No known examples
Content or rule-based analysis	Descriptive information about alternatives and user interests. Match alternative profile to user profile.	Information items – information retrieval or filtering	Syskill & Webert [14] [http://www.ics.uci.edu/~pazzani/] NewsWeeder [10]
Collaborative systems – consult peers	The judgements or opinions of a like group of peers	Choices involving tastes; alerting and filtering systems	Phoaks [22] [http://www.phoaks.com/phoaks/]
Rating Agencies - consult domain expert(s)	The judgement or opinion of one or more experts in the choice domain	Problems in an organizational learning application; issues of high importance	Referral Web [8] Argus Clearing House [http://www.clearinghouse.net/]

Table 1: Recommendation System Approaches

Utility Based Techniques: In this approach, the expertise (on what the user requires) is assumed to reside only in the user. The system interacts with the user to construct an explicit representation of the user's utility function in terms of the features (attributes) of the choices. Utility estimation techniques are suitable for choices between extremely

valuable items such as cars and houses for which attribute values such as price are readily available, but for which the ultimate choice must be made by the user. The Prefcalc system [11]³, for example, requests users to rank a subset (five or six items) from a much larger set of choices and uses the rankings to construct an additive piece-wise linear utility function for the user. Given values for the attributes of each possible choice (such as the price, size, speed and fuel consumption in the case of automobiles), it is then a simple matter to rank all the choices - even those for which the user had no prior knowledge. In a similar vein, users could interactively and implicitly indicate their utility for the various choice possibilities as in multi-criterion decision making techniques [25]. While a scoring approach based on utility theory is possible and even desirable in some situations, to our knowledge, no existing Web-based recommendation systems attempt to estimate user utility functions in this classical, decision theory sense.

Content-Based Systems: In this approach, the knowledge required to rank alternatives is embedded in a combination of descriptive information about the alternatives themselves (e.g. document terms or attribute values) plus an explicit user profile (e.g. keywords indicating subject interest, or a past history of previous choices by the user.) The content-based (or rule-based) approach has its roots in information retrieval (IR) [2]. The techniques used for full-text searches of Web-based documents are similar to those used in IR. The key to improved performance of the content-based approach over traditional IR approaches is the use of user profiles containing information on user preferences and tastes. These may be explicitly elicited through questionnaires or built over time by observing and tracking users as they interact with the Web. Relevance feedback, an important component of content-based approaches, is used to update user profiles. One or more of the above-mentioned IR methods may be used for weighting words to represent documents and text. Similarly several different methods exist for updating user profiles.

"Syskill & Webert" is the name of a software agent that uses a content-based approach to recommender systems [14]. Syskill & Webert learns the user's interests and preferences and then uses a LYCOS (Web-based search engine) query to retrieve Web pages that match the user's profile. The user evaluates the retrieved Web pages, which are usually related to a narrow subject domain, and these preferences are then stored in the user's profile, which is updated as more evaluations are made. A simple Bayesian classifier is used to determine the probability that the user would like a Web page. The system uses a boolean

³ Prefcalc is a stand-alone MCDM application (not web-based.)

vector space model for learning 'features' (document terms) that help discriminate between interesting sites and uninteresting ones.

Content-based approaches build upon IR models and consequently inherit many of their limitations. In particular, given the diversity of resources on the Web, not all Web-based documents are amenable to proper representation using traditional IR indexing techniques, which are best suited for text-based documents. Secondly, content-based approaches rely heavily on prior feedback from users. This results in 'over-specialization' i.e., the system performs best in restricted domains that the users have evaluated in the past [2]. Thirdly, obtaining proper feedback from users for retrieved items is time-consuming and frustrating - given their limited knowledge and experience, users cannot determine what documents best satisfy their needs without actually browsing through each and every one of them. More importantly, the user might not be in a position to satisfactorily evaluate the quality of the items retrieved. 'Collaborative' recommender systems overcome some of the limitations of content-based systems, but have some weaknesses of their own.

Collaborative Systems: In collaborative approaches, the opinions of people that have similar interests and tastes to the user provide the basis for the recommendations. This is the basic philosophy underlying recommender systems (narrow definition) and collaborative filtering applications. "Rather than compute the similarity of the items, the collaborative approach computes the similarity of users" [2]. In collaborative systems, the user profile does not represent the user's preferences over the various choices that are to be recommended. Rather, it helps determine the user's similarity to other users of the recommender system. The recommendations are then based on the choices made by users with similar profiles. For example, a user might simply indicate a set of keywords of interest and will be provided rankings by a subset of users with a similar profile. In this approach the software agent finds other users whose preferences are similar to those specified by the user and then recommends items that they liked. Clusters of users with similar preferences are identified based on the correlation of their earlier evaluations of items in the choice domain. This, however, requires that the user have rated the same items for comparison. Scores for unseen items are predicted based on a combination of scores known from the nearest neighbors in a cluster [2].

PHOAKS (People Helping One Another Know Stuff) is a collaborative recommender system that mines Usenet News groups for mention of Web pages (URLs), which on passing a number of tests are then classified as recommendations [22]. Some of the search, classification and filtering techniques

common to IR systems are used. URLs that are cross-posted to a number of newsgroups and have accompanying text that are suggestive of an advertisement or promotion are automatically discarded. URLs recommended by a large number of users were found to be of higher quality than those with fewer recommendations. Additional information from FAQ (Frequently Asked Questions) databases was used to measure and improve the quality of the recommendations. One of the limitations of the PHOAKS system is that it does not distinguish between different evaluations, thus assigning them equal weights, irrespective of their credibility.

Collaborative approaches can recommend a wide variety of items, not just those evaluated by the user. One of the most significant advantages of the collaborative approach is how little an individual user has to contribute to be able to retrieve relevant documents since the approach relies less on the user's own evaluations and more in the evaluations of users with similar tastes. However, the success of such a system depends on the availability of a critical mass of users with similar profiles, on the willingness of the users to contribute evaluations, and the credibility of their evaluations. As some users who benefit from such systems have no incentive to contribute evaluations, there may be problems of free riding. Market-mechanisms and pricing schemes for evaluations have been suggested to overcome some of these problems [16].

Third Party Expertise: In the fourth approach to producing recommendations, the system facilitates consultation with a domain expert. This approach might be indicated for problems and issues of major consequence. The referral web [8], in which users are linked into communities around domain specialists, takes this approach.

The Argus Clearinghouse (<http://www.clearinghouse.net/>) serves as a clearinghouse for topical guides that identify, describe and evaluate Internet-based information sources. The topical guides are rated based on five criteria: (i) Level of resource description (content, currency, access, technical performance etc.); (ii) Level of resource evaluation (subject quality indicators, information on authors, document layouts, graphics etc.); (iii) Guide design (images, layout, navigational aids etc.); (iv) Guide organizational schemes; and (v) Guide meta-information. Exceptionally good guides are also given a "Digital Librarian's Award". Rating agencies (or Clearinghouses) can thus serve as a credible source of recommendations in different domains.

Rating agencies or third party experts are useful in situations where special expertise is required in order to make the recommendations. A disadvantage is that the recommendations produced by such systems are not usually personalized to the user.

Hybrid Systems – Most existing recommendation systems take a hybrid approach involving some combination of content-based analysis and collaborative filtering. The Barnes and Noble implementation of the Firefly software mentioned earlier uses a combination of rule-based and collaborative filtering.

FAB – The FAB system, a part of the Stanford University digital library project, is a hybrid recommendation system [2]. It is comprised of three main components – *collection agents* that find pages on a specific topic; *selection agents* that find Web pages for a specific user; and *the central router*. Users receive items both when they score highly against their own profile as well as when they are rated highly by a user with a similar profile. While the collection agent's profile represents its current topic, the selection agent's profile represents a single user's interests based on his/her evaluation. The central router forwards Web pages from the collection agent to users based on their profiles, which are stored in their personal selection agents. Collection agents specialize in narrow domains and adapt to a dynamically changing population of users. The FAB system uses Web based full-text search engines to retrieve pages that match the user's profile. The success of the system depends largely on how accurately the users' profiles reflect their true preferences.

Hybrid systems overcome some of the limitations of content-based systems. By utilizing group feedback we potentially require fewer cycles to achieve the same level of personalization [2]. In addition, individuals gain because the group can appraise more items and will usually provide a different viewpoint.

The application domains in Table 1 are suggestive only. Each technique could be applied in multiple domains. This is suggested by Table 2, which arranges the information described above around some typical problem domains. (The third row in the table lists the alternative approaches in each problem domain).

MANAGERIAL IMPLICATIONS

From a firm's viewpoint, recommendation systems can be used in two ways – as a marketing tool to positively influence consumer perceptions and preferences, and/or as a tool for knowledge management and organizational learning. Table 2 shows sample problem domains in these two areas of concern. The first and second rows in Table 2 list the choice domain and characterize the type of problem to be solved by the recommendation system, the third row suggests alternative sources of recommendations, while the last row describes the value that the system provides its users.

As is to be expected, the design and deployment of recommendation systems should differ to reflect the concerns of the different problem domains. Columns 2 and 3 illustrate two contrasting decision situations in e-commerce - high value, one-off decision making in which personal preferences are of overriding importance and low value decision making where matters of opinion and taste are paramount. The level of trust, risk, validity and performance expected in these two situations is vastly different. For example, as suggested in the third row of the table, the locus of decision making for high value decisions should probably reside in the user, while in low value decision making situations it seems feasible to rely on more automated approaches.

	Electronic Commerce		Knowledge Management
What is recommended?	Durables – known attribute values (such as price)	Consumables- matters of opinion	Continuous flow of knowable product e.g. information
Typical choice domains	Cars, houses, vacations	Art, movies, books, CDs	News feeds, Usenets, problem databases, organizational learning
Problem type	Optimal selection from a relatively few big ticket items	Selection from many items of low value.	Filtering; prevention of information overload; directing attention to important issues
Who is the recommender?	(1) Utility-based (User + decision aid) (2) Rating Agencies	(1) Content-based (Rules + personal profile) (2) Collaborative (Peer group) (3) Rating Agencies	(1) Content-based (Rules + personal profile) (2) Collaborative (Peer group) (3) Experts
Value added by recommendation System	Information gathering and presentation; one-to-one matching to personal preferences	Information gathering and presentation; one-to-one matching to personal preferences; save time and cost of sampling; align with opinions of others	Information filtering, evaluation and presentation; alerting to new issues and ideas; match to personal information needs; save time by reading only items of direct interest

Table 2: Problem Domains, Alternative Approaches and Value Added

As organizations are overwhelmed by the information explosion and as they become more knowledge intensive, organizational learning and knowledge management within the organization becomes important. Recommendation systems, using some of the techniques discussed earlier, can help firms in this area. Column 4 in Table 2 lists

information filtering applications in which the chief concern is to bring only interesting or new (from the recipients point-of- view) items to the attention of employees and to eliminate uninteresting or unimportant items.

ISSUES SURROUNDING THE USE OF RECOMMENDATION SYSTEMS

Space does not permit a review of the many interesting issues and problems involved in successfully implementing a recommendation system. Obviously, these issues span the gamut from the development of advanced software technologies to the care and nurturing of customers. In the case of recommender systems in particular, the developers must be concerned with how they can develop a community of people who will actively contribute to the common good. In this section, we briefly discuss only two questions: Why would a user accept the recommendations of a recommendation system? And, how can the value of such systems be measured?

The major factors associated with user acceptance and use of recommendation systems are shown in Table 3. The issues of trust, validity, privacy, and risk are necessary preconditions for acceptance of recommendation systems, but ultimately, their success, will depend on the performance factor - the utility they provide their users.

Trust:	Authenticity of source - is the recommender who I think it is? Does the recommender provide an unbiased recommendation?
Validity:	Does the system have the correct information/expertise? Does the system understand my requirements and wishes? How does the system make its recommendations?
Privacy:	Is my personal profile kept secret? Are my inquiries and purchases private?
Risk:	What redress do I have if I act on bad recommendation? (e.g. Money back if not satisfied?)
Performance:	What value do I derive from the recommendations - do I save money or time?

Table 3: Factors Influencing the Acceptance of a Recommendation System

Trust - Possible bias in the recommendations due to self-interest of either the owners or users of such systems is a major issue, particularly for collaborative recommender

systems. "If anyone can provide recommendations, content owners may generate mountains of positive recommendations for their own materials and negative recommendations for their competitors," [16].

Validity - The question here is whether the system has the information and expertise to be helpful. Just as we associate more or less validity to the pronouncements of wine and food critics or human experts in any other walk of life, so too will computer recommendation systems acquire reputations for the relative usefulness or otherwise of their recommendations. The second question under this heading in Table 3 refers to the ability of the recommendation system to adapt its response to the needs of individual users. The third question emphasizes the need for recommendation systems to have accurate and reliable inference mechanisms - an active area of current research.

Privacy - Most recommendation systems rely upon the evaluations and ratings provided by different users to recommend items. Sharing the opinions of individuals with the society at large raises concerns relating to privacy. While anonymity and the use of pseudonyms provide a partial solution to this problem, several alternatives have been proposed. For example, the Open Personalization Standard (OPS) also known as the Open Profiling Standard is yet to be agreed upon by industry participants [6]. Under OPS, users could choose to store personal information, hobbies and interests on their PC hard drives and then decide whether to disclose that information to a particular web site.

Risk - The risk assumed by users of recommendation systems varies with the choice domain. For users of filtered Usenet information the risks are probably inconsequential, but for users of recommendation systems for large budget items or stock advisory services, the risks of acting upon computer-generated recommendations will be quite large. In normal commerce, both competition and legal considerations have shaped the way risks are shared between sellers and consumers. We suspect that a proper understanding of the risks for both providers and users of recommendation services will be similarly shaped over a period of use.

Performance - The market for recommendation services will ultimately decide which ones survive and prosper. Performance will improve over time as the technology becomes more sophisticated and the surviving intermediaries reach and surpass critical mass. Size is especially important for recommender systems, which depend on

evaluations and feedback from users and hence, require a critical mass of recommenders and recommendations to be able to perform effectively.

This leads us to the interesting question of how the performance of recommendation systems can be evaluated. One possibility is to use the notions of recall and precision that were defined above. More generally, we stated in the introduction that the role of a recommendation system is to increase the probability that its users will be satisfied by the choices that they make as a result of using the system. From a similar viewpoint, we have the notion of "predictive utility" [9]. In other words, recommendation systems must be able to predict which items in its domain of choice will most satisfy each user. A simple cost-benefit analysis approach based on [9] illustrates the need for a closer look at performance measures for recommendation systems. Table 4 shows the possible outcomes for a recommendation system that makes binary predictions as to whether an item will be useful to the user (good) or not useful (bad). Illustrative costs and benefits are provided for two different hypothetical recommendation systems: one that rates movies and a second that filters news items for financial analysts. To illustrate the kind of reasoning that might determine the values in each cell, the benefit of the correct prediction of a good movie (one that the user will enjoy) is listed as "medium", while the cost of a false positive (the prediction that the user will enjoy the movie, when he/she does not), is the cost of buying a ticket and wasting an evening - which has been judged to be "high." As another example, the value of a hit and the cost (opportunity cost) of a miss of a pertinent news item to a financial analyst are both rated as "high" in the table.

	Predict "Good item"	Predict "Bad item"
Good Item	HIT Movie: Medium benefit News item: High benefit	MISS Movie: Low cost News item: High cost
Bad Item	FALSE POSITIVE Movie: High cost News item: Medium cost	CORRECT REJECTION Movie: Medium benefit News item: Low benefit

Table 4: Cost-Benefit Analysis for Recommendation Systems

As discussed earlier, recall measures the number of relevant items ("hits") divided by the total number of items in all four quadrants, while precision records the number of hits

divided by the total number of hits plus false positives. But such ratios do little to reflect the costs and benefits of the various possible outcomes and can therefore, shed little light on whether a recommendation system is beneficial to an individual or an organization. A complete cost-benefit analysis, even in such simple applications, is quite difficult because it depends on a correct assessment of the utility of the various outcomes for each particular user as well as the probabilities of a given recommendation ending up in each of the four cells⁴. Nevertheless, a consideration of the issues raised by even an approximate cost-benefit analysis should provide useful guidance for the designers of recommendation systems. Further research on the costs and benefits of recommendation systems is urgently required.

SUMMARY AND CONCLUSIONS

To cope with the ever-increasing complexities of the information economy, we suggested in the introduction that research and development of a range of "coping technologies" is needed. In fact, we believe that we will need to develop support environments that will, among other things, monitor the environment, direct our attention to what is urgent, relevant and important, increase our understanding of the world and help us with our decision making. Recommendation systems are an important new technology that may help us in the decision-making aspect of our lives. We defined recommendation systems broadly as systems that provide recommendations by any means, reserving the term recommender systems for the important sub-class of recommendation systems that make recommendations based on the opinions of other people.

The explosive growth of the Internet and the World Wide Web has resulted in the proliferation of information of uncertain quality in widely varying formats, which has greatly reduced the effectiveness of traditional IR approaches. More importantly, the needs of electronic commerce extend far beyond those of information dissemination and retrieval. In this new world, we see a huge potential market for systems that collect and analyze user tastes and needs on the one hand, and signal the quality of products and services on the other hand. Throughout the chapter, we discussed a number of examples of expert-based, content-based, collaborative and hybrid recommendation

⁴ A closed form solution for this problem was first proposed in the IR context by [24].

systems. Some of these are research systems; others are already commercial successes.

The framework for recommendation systems clarifies the differences with IR systems and the general functions provided by recommendation systems. We also identified four general classes of recommendation systems based on the source of expertise underlying the recommendation: utility-based systems, content-based systems, collaborative systems, and expert consultation. Finally, we briefly discussed the issues of user acceptance of recommendation systems and the need for research on performance measures.

This brief survey and classification of recommendation systems has barely scratched the surface of a whole new technology of assisted decision making that we feel will grow in importance as electronic media become the main vehicles for human communication, education and commerce. As we argued in the introduction, recommendation systems are a response to the need for individuals and organizations to manage the overwhelming flood of new information, products and services, and to make more and more decisions under increasing time pressure. For these reasons, we believe that recommendation systems are inevitable and that they will be part of a thriving new industry of web-based intermediaries.

The social consequences of such systems are enormous. If recommendation systems fulfill their promise, they could make markets more efficient by providing consistent, valid and credible quality signals thereby reducing or eliminating the hassle associated with search and investigation of a wide range of goods and services. Recommendation systems also provide opportunities for one-to-one marketing, which, by catering to individual tastes, should increase overall welfare.

On the other hand, there are several dangers. First, individuals and organizations could become overly dependent on such systems creating opportunities for unscrupulous companies, or even government agencies, to manipulate user tastes and decision making. Second, it is not at all certain whether widespread use of recommendation systems will lead to more diversity and freedom of expression or to the opposite - an economic and social system in which individual tastes and needs are satisfied on the margin, but in which overall cultural and intellectual directions are dictated by the majority. Civil liberty groups, for example, are concerned by the power of filtering systems to limit free speech on the Internet [13]. In our opinion, the future

economic and social consequences of these new technologies are a matter for urgent research and debate in academic, industry and governmental regulatory agencies.

REFERENCES

- [1] Alter, S. "A Taxonomy of Decision Support Systems," *Sloan Management Review*, 1977, 19, 39-56.
- [2] Balabanovic, M., and Shoham, Y., "Fab: Content-Based, Collaborative Recommendation", *Communications of the ACM*, March 1997, 40, 3, 66-72.
- [3] Belkin, N. J., and Croft, B. W., "Information Filtering and Information Retrieval: Two Sides of the Same Coin", *Communications of the ACM*, December 1992, 35, 12, 29-38.
- [4] Belkin, N. J., Oddy, R. N., and Brooks, H. M., "ASK for Information Retrieval. Part I. Background and Theory", *Journal of Documentation*, 1982, 38, 61-71.
- [5] Gorry, G. A. and M. S. Scott Morton, "A Framework for Management Information Systems", *Sloan Management Review*, Fall 1971, 13, 1, 55-71.
- [6] Infoworld, "Web Marketing Gets Personal", January 12, 1998, 20, 2, 93-94.
- [7] Kambil, A., "Doing Business in the Wired World", *IEEE Computer*, May 1997, 56-61.
- [8] Kautz, H., Selman, B., and Shah, M., "Referral Web: Combining Social Networks and Collaborative Filtering", *Communications of the ACM*, March 1997, 40, 3, 63-65.
- [9] Konstan, J.A., Miller, B.N., Maltz, D., Herlocker, J.L., Le, G.R., and Riedl, J., "GroupLens: Applying Collaborative Filtering to Usenet News", *Communications of the ACM*, March 1997, 40, 3, 77-87.
- [10] Lang, K., "Newsweeder: Learning to Filter Netnews", *Proceedings of the 12th International Conference on Machine Learning* (Tahoe City, Calif.), 1995.
- [11] Jacquet-Lagrèze, E. and J. Siskos, "Assessing A Set Of Additive Utility Functions For Multicriteria Decision-Making: The UTA Method," *European Journal of Operational Research*, 1982, 10, 151-164.
- [12] Malone, T.W., Yates, J., and Benjamin, R.I., "Electronic Markets and Electronic Hierarchies", *Communications of the ACM*, June 1987, 30, 6, 484-497.
- [13] Marmon, Amy, "Technology to Let Engineers Filter the Web and Judge Content," *New York Times*, Monday 1998.
- [14] Pazzani, M., Muramatsu, J., and Billsus, D., "Syskill & Webert: Identifying Interesting Web Sites", <http://www.ics.uci.edu/~pazzani/>
- [15] Pine, B., Joseph III, Don Peppers and Martha Rogers, "Do You Want to Keep Your Customers Forever?" *Harvard Business Review*, March-April 1995, 103-114.
- [16] Resnick, P., and Varian, H.R., "Recommender Systems", *Communications of the ACM*, March 1997, 40, 3, 56-58.
- [17] Robertson, S.E., "Theories and Models in Information Retrieval", *Journal of Documentation*, 1977, 33, 126-148.

- [18] Salton, G., and Buckley, C., "Term-weighting approaches in automatic text retrieval", *Information Processing and Management*, 1988, 24, 513-523.
- [19] Saracevic, T., "Relevance: A Review of a Framework for Thinking on the Notion of Information Science", *Journal of the American Society for Information Science*, 1975, 26, 321-343.
- [20] Stein, E., and Zwass, V., "Actualizing Organizational Memory with Information Systems", *Information Systems Research*, June 1995, 6, 2, 85-117.
- [21] Stohr, E.A., and Viswanathan, S., "Web-based agents for Electronic Commerce", Unpublished Manuscript, Center for Research in Information Systems, Stern School of Business, New York University, 1998.
- [22] Terveen, L., Hill, W., Amento, B., McDonald, D., and Creter, J., "PHOAKS: A System for Sharing Recommendations", *Communications of the ACM*, March 1997, 40, 3, 59-62.
- [23] Tuzhilin, Alexander, "The E-Butler Service, or Has the Age of Electronic Personal Decision Making Assistants Arrived," Center for Research in Information Systems, Stern School of Business, New York University, 1998.
- [24] Verhoeff, J., Goffman, W., and J. Belzer, "Inefficiency of the Use of Boolean Functions for Information Retrieval," *Communications of the Association for Computing Machinery*, 1961, 4, 557-558.
- [25] Zionts, S. and J. Wallenius, "An Interactive Method for Solving the Multiple Criteria Problem," *Management Science*, 22, 4, 1978.