

Social Recommendations Systems: Leveraging the Power of Social Networks in Generating Recommendations

Ofer Arazy¹
School of Business
University of Alberta
Edmonton, Alberta, T6G 2R6, CANADA
ofer.arazy@ualberta.ca

Nanda Kumar
Computer Information Systems Department
Zicklin School of Business
Baruch College, City University of New York
New York, NY 10011, USA
Nanda_Kumar@baruch.cuny.edu

Bracha Shapira
Department of Information Systems Engineering
Faculty of Engineering Sciences
Ben-Gurion University of the Negev
Beer-Sheva, ISRAEL 84105
bshapira@bgumail.bgu.ac.il

Abstract

Recommendation systems, and specifically Social Filtering (SF) systems, play a significant role in reducing information overload and providing users with information relevant to their specific interest. For over a decade now, the ad-hoc standard in social filtering employed an approach, where recommendations were generated by computing “shared interests” based on users’ preferences for items. The rapid growth in online social networks presents an opportunity for a new social filtering approach. The main thrust of our work is in identifying the relevant relationship characteristics among participants who know each other and use these characteristics to improve the quality of the recommendations generated. This paper develops a model that incorporates users’

¹ Authors’ names are presented in alphabetical order.

explicit perception of sources' trustworthiness to improve the quality of the recommendations and proposes an experimental design to test the model. We plan to conduct the experiment in November 2005 and will have results ready for presentation well in time for the conference.

1. Introduction

Information overload impedes organizations' performance. The rapid growth of available information online, fueled by the rapid adoption of the internet and the web, is making access to relevant information analogous to that of finding a needle in a haystack. This problem is accentuated in today's markets, where information and knowledge play a critical role in firms' competitive positioning.

Information Filtering (IF) systems play a significant role in reducing information overload and provide users with information relevant to their specific interests. IF systems are now an integral part of firms' information architecture, serving both customers (e.g. Amazon's book recommendations) and internal employees. Over the last decade, much research has been done specifically on social/collaborative IF systems by both academia and industry. Social Filtering (SF) generates quality recommendations by identifying sources whose recommendations could be trusted and using the ratings (i.e. quality evaluations) of these trusted sources (Shardanand and Maes, 1995), as illustrated in Figure 1.

One specific operationalization of the SF model, where a source's trustworthiness is estimated based on the degree to which receiver and source share interests (Shardanand and Maes, 1995), has gained dominance over the last decade. In this paper, we use the term 'Shared Interests' to refer to this approach. This seminal operationalization has

served as the basis over which many variations have been built. Most of the current SF systems typically use the Shared-Interests approach (e.g. Amazon's recommendation system) as it is easy to automatically establish the relations between individuals. Users' ratings could be gathered either implicitly (e.g. based on a purchase or other recorded transaction) or explicitly (user ratings) without requiring much effort. This data is then used to identify similar individuals, thus generating relationships automatically. The limitations of the Shared-Interests approach stem from the fact that similarity of interests may not be the best proxy for sources' trustworthiness (the bedrock of the original operationalization), and may not produce the best possible recommendations.

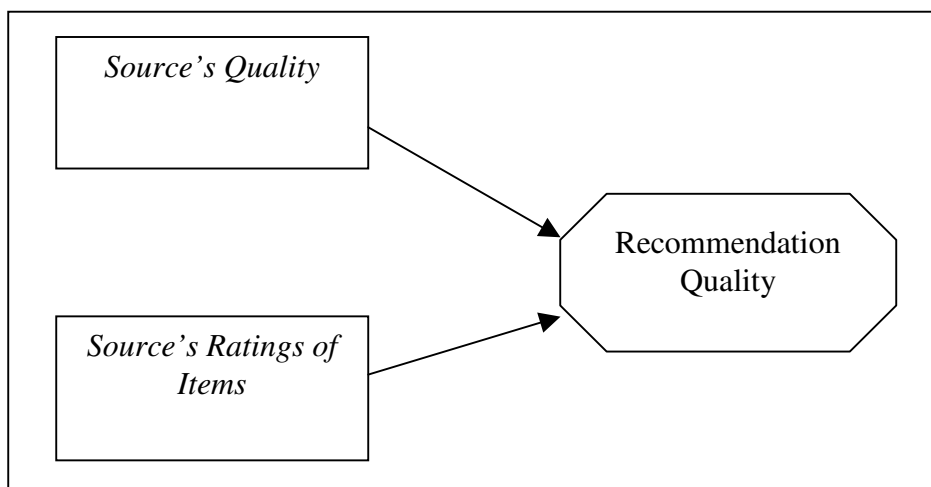


Figure 1: a conceptual diagram of the Social Filtering model

However, the emergence of social networks on the internet in recent years - professional (e.g. LinkedIn) and recreational (e.g. MySpace) communities - provide new ways of incorporating sources' trustworthiness in generating recommendations. In online social networks, users establish relations with other users, and these relations are used for communicating and sharing information. To the extent that these relationships indicate common trust among users, they could potentially be utilized by SF systems.

Specifically, SF systems could employ explicit relations of online communities as an indicator of source's trustworthiness in addition to using Shared Interests (or possibly substitute the Shared-Interests measure). Thus, the coupling of online social networks with recommendations systems has the potential to enhance filtering system's quality. The growing interest in this area in academia as well as the emergence of social information access systems (such as Yahoo's MyWeb) illustrates the promise of this approach.

Notwithstanding the potential of social networks for SF, some critical questions remain unanswered. It is not clear if the links in online communities indicate trust relations, and whether these links have the necessary characteristics required for ensuring quality recommendations in filtering systems. For instance, the links in Professional Communities, such as LinkedIn, may entail practical or opportunistic relations, while links in dating sites most likely entail romantic relations.

This work presents the first deliverable out of a larger project that looks to establish theoretical foundations for collaborative filtering systems in the social networking context. The project hopes to demonstrate that theory-driven design, along the lines suggested by Design Science research methodology (Hevner et al. 2004), could result in improved SF system performance. The objective of the paper is consistent with the growing emphasis in the MIS discipline on the importance of directly linking the study of perceptions and intentions to the design of IT artifacts (Benbasat & Zmud 2003). Specifically, we investigate the type of relationships between recommendation receiver and source that would yield the best possible recommendation quality. This initial work has two broad objectives: (1) to identify the relationship factors that are expected to

impact recommendation quality, and (2) to design an empirical experiment that would test the impact of the identified factors on SF system performance. The work presented here is in early stages and preliminary results of the experimental study are expected to be available by the Design Science conference date.

This paper is organized as follows: Section 2 describes the prior works on the recommendation process, filtering systems, and online communities; Section 3 introduces our proposed model; Section 4 describes an experimental design aimed at testing the proposed model; and Section 5 discusses the expected contributions and extensions of this study.

2. Prior Works

Information Filtering systems have been investigated for over three decades and have been applied in various contexts – from filtering-out irrelevant e-mail to the recommendations of products to be purchased. For surveys of the field, please refer to Hanani et. al, (2001), and Oard (1997).

Early IF systems, usually deployed in restricted organizational settings, required the users to explicitly define their topics of interest by choosing relevant keywords or categories in order to generate their profiles². The filtering process consisted of a comparison between the users' profiles and incoming data streams that were then classified as relevant or non-relevant based on the users' preferences. The main limitation of this approach is that it required significant user effort in defining the initial profile and in continually updating the profile to ensure that it reflects user's evolving needs. Another

² Models and techniques for generating – either manually or automatically – user profiles have been studied in the area of User Modeling. For a survey of this field please refer to (Kobsa 2001).

problem with this approach was that users found it difficult to define their preferences as a set of keywords (Boger et. al., 2001). It was found that the resulting profiles were not very useful for filtering³. As a result, traditional filtering implementations remained restricted to organizational settings, and did not gain popularity on the internet.

Users / Items	Relationship Data	Data for Predictions		
		Item 1	Item 2	Item 3
User A	Trusted information sources	0.4	-	0.7
User B		-	0.2	0.6
User C		0.5	0.3	-
User D		0.6	-	-
User E		-	0.8	0.2

Data used for predicting a recommendation for *UserD* on *Item5*

Figure 2: the social filtering model.

In the mid 90s, social (or collaborative) filtering (SF) emerged as an alternative filtering approach (Goldberg et al. 1992, Shardanand & Maes 1995). In SF, prediction of the relevancy of an information item and its recommendation to a user is viewed as a social process. The key to delivering relevant information is in the identification of individuals whose recommendations could be trusted, and in leveraging the experiences of these trusted people to generate accurate recommendations (Shardanand & Maes

³ An alternative technique for constructing user profiles is the implicit acquisition of users' preferences. Users' preferences could be automatically extracted by monitoring their browsing behavior and automatically (i.e. without disrupting the normal pattern of users' work) inferring users' relevance judgments. For example, significant correlation was found between the time spent on an information item and the relevancy of the information for the user (Morita & Shinoda, 1994; Claypool et. al., 2001). Alternative indicators of relevance, such as mouse clicks, printing, or book-marking and their combinations, were also examined. Notwithstanding these encouraging results, the effectiveness of automatic approaches in predicting users' preferences remains questionable, and their accuracy is still inferior to that of explicit acquisition (Kelly & Teevan, 2003).

1995). SF consists of (1) recording users' ratings (relevance evaluations) of items they are exposed to, (2) identifying trusted sources of information for every user; and (3) predicting the relevant items for users based on items rated favorably by the trusted sources. The SF conceptual model is shown in Figure 1, and its architecture is illustrated in Figure 2.

Users / Items	Training Data			Data for Predictions		
	Item 1	Item 2	Item 3	Item 4	Item 5	Item 6
User A	0.5	-	0.9	0.4	-	0.7
User B	0.2	0.8	-	-	0.2	0.6
User C	-	0.4	-	0.5	0.3	-
User D	0.3	0.1	-	0.6	-	-
User E	-	0.5	0.3	-	0.8	0.2

Data used for predicting a recommendation for *UserD* on *Item5*

Data used for determining *Shared Interests*

Figure 3: the Shared Interests SF approach, with a set of training data, used for calculating user-user similarities, and the prediction data.

The earliest work to propose SF (Goldberg et al. 1992) suggested that relationships between the receivers of recommendations and their respected trusted sources be defined explicitly. However, this SF operationalization still required significant user effort, as users had to manually define the set of trusted sources and rate information items for their relevancy. An alternative operationalization of SF was proposed three years later (Shardanand & Maes 1995), where receiver-source relationships were extracted automatically, based on similarities between recorded logs of users' ratings. This SF approach (referred to as the 'Shared Interests' approach) became

popular on the internet because it reduced user intervention. Recent systems attempt to combine this approach with implicit acquisition of users' interests to eliminate the need for any user effort. Thus, in a fully-automatic SF, users' ratings are implicitly gathered (e.g. based on a purchase or other recorded transaction), and the recorded history of users' ratings is used to identify similar individuals, and thus relationships are also elicited automatically (see Figure 3).

Another advantage of the Shared Interests approach was the large number of recommendations sources that were made available for each user. Access to the ratings of these large number of sources (with shared interests were previously unknown to the user and may remain unknown) is an important factor in ensuring recommendations quality. In recent years, the Shared Interests SF approach has attracted significant attention, and research in this area has provided enhancements along various dimensions such as automatic elicitation of accurate user feedback, algorithms for measuring users' similarities, and improving prediction methods, resulting in better system effectiveness over alternative filtering approaches (e.g., Kelly and Teevan 2003). In industry, the 'Shared-Interests' approach has proved extremely popular and most online filtering systems (e.g. Amazon's recommendation system⁴) identify trusted sources based on users' similarity of interests.

The limitations of the Shared-Interests approach stem the fact that it captures only one dimension of receiver-source relationship. In fact, this approach is 'social' only in a weak sense, since users that were found similar may not even know one another, let alone collaborate or engage in social interaction. We believe that the relationship characteristics

⁴ Amazon.com uses a variation of this approach known as item-to-item collaborative filtering to solve sparsity problem (it is easier to find items that are rated by many users than users that rate the same items).

that are necessary for producing quality recommendations could be much richer than the automatically-extracted similarity of interests. The main thrust of our work is in identifying the relevant relationship characteristics among participants who know each other and use these characteristics to improve the quality of the recommendations generated by the previous generation SF systems that depend solely on the shared interests among unknown participants.

Management literature has studied decision-making processes extensively, but only recently has begun to investigate the recommendation and advice processes. It is generally accepted that recommendation quality depends on source's trustworthiness (Hardin 1993; McKnight et al. 2002; Levin et al. 2004; Guha et al. 2004). While numerous models of trust has been proposed in recent years, we use the conceptualization proposed by Mayer et al. (1995), where 'Trustworthiness' includes three dimensions: *Integrity*, *Benevolence*, and *Competence*. Management literature suggests that the quality of source's recommendations depends on his Integrity, Benevolence, and Competence. Integrity implies truthfulness, honesty, and ethical conduct; Benevolence implies liking, good-will, and affection; and Competence denotes expertise, knowledge, and skills, and is usually context-specific.

Hence, we believe that SF techniques that will rely on these types of receiver-source trust relations will yield enhanced prediction accuracy. It is important to note that the studies reviewed above were usually conducted in a corporate setting, with relatively few subjects, where most of the social relations were restricted to the workplace. These settings are very different from the settings in which large-scale social networks operate, i.e. a large user-base and leisure-related social relations that characterize the internet

environment. Thus, it is not clear whether findings from these studies are applicable to the design of SF systems on the internet.

The emergence of online social networks presents a great opportunity for filtering systems, namely in exploiting the relationships embedded in a social network to detect trusted sources for generating recommendations. This approach truly deserves the name ‘social filtering’, as the receiver-source links captured are meaningful relations between individuals (over and above the Shared Interests links). In industry, Epinions.com is one of the most notable example of a commercial social filtering system, where users’ trust relationships form a “web of trust”, which is then utilized to filter product recommendations. In a closely related field - web search - we are now witnessing the emergence of new social systems, such as del.icio.us and Yahoo’s MyWeb (still in Beta version).

In academic research, since the early work of Goldberg et al. (1992), very few works addressed the issue of explicit social relations for establishing trusted sources in filtering systems. In a few recent formulations of trust, it has been suggested that explicit trust relations could be utilized by search and filtering systems (Abdul-Rahman & Hails 1999; Guha et al. 2004). Gnasa et al. (2004) propose a search architecture that includes social relations (although these are restricted to shared interests), and have developed a prototype (Iskodor) to illustrate their ideas. Freyne and Smith (2004) introduce the I-Spy search engine, which employs un-named explicit social relations (and the recorded browsing experiences of related individuals) to enhance search performance within defined communities that share interests. The studies reviewed above are works-in-

progress that demonstrate the rising interest in SF systems with explicit receiver-source relationships.

3. Proposed Model

The literature suggests that explicit receiver-source trust relations are expected to enhance SF performance, and thus social networks could potentially be exploited by SF systems to enhance prediction accuracy. We propose an architectural model for social filtering systems, in which sources are qualified based on their trustworthiness, as well as on the degree to which they share interests with source. Figure 4 below presents the proposed model.

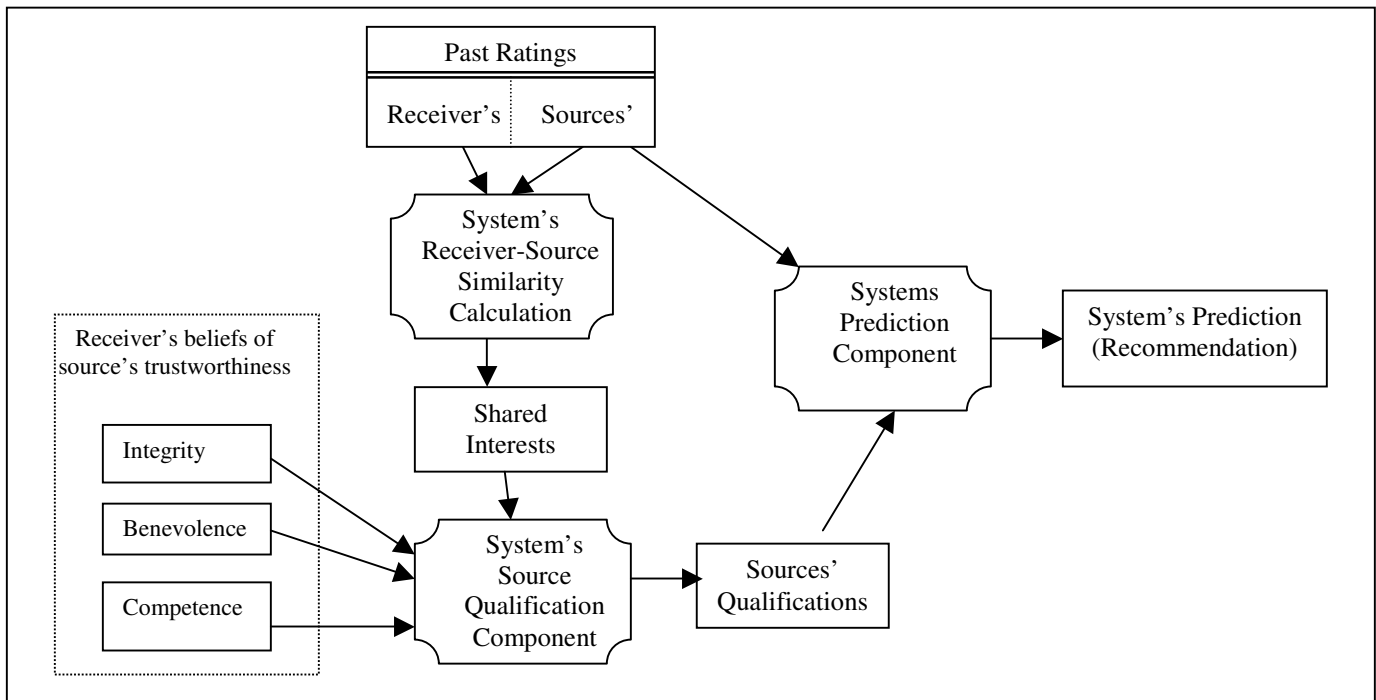


Figure 4: the proposed model, linking source's trust, source quality and SF accuracy.

As illustrated above, the prediction component of the SF system in the proposed model receives two types of inputs about each source. The first is the quality of the

source computed by the System's Source Qualification Component by taking into account the receiver's perception of the source's trustiness as well by the shared interest between the receiver and source. The receiver's perception of trustworthiness is computed based on the three trust dimensions: - *Integrity, Benevolence, and Competence*. The effect of each of these dimensions on the receiver's perception is to be determined empirically by the experiments that we plan to conduct. Also the relative importance of shared interest and the source's trustworthiness on the source's quality is to be empirically determined. The other input to the prediction component is the source's ratings on relevant items. The output of the component is a recommendation of information items to users that considers the trust between sources and receivers. The prediction of an item to a user is thus based on computing the quality of each of the sources and considering their rating for the item to be predicted relative to their quality. A rating provided by a "good source" should have more effect on the recommendation than inferior sources. The function that computes the quality (Weight) of the source S_k for user i is defined as follows:

$$W_{Sk}(i) = \delta T(i, Sk) + \lambda SH(i, Sk)$$

Where:

W =Weight, T =trust, and SH = shared interest

δ, λ - denote the relative importance of the trust and shared interests on the quality of the source, they sum up to 1.

The trust is composed from three dimension:

$$T(i, Sk) = \alpha I(i, Sk) + \beta B(i, Sk) + \gamma C(i, Sk)$$

Where:

$T=Trust, I=Integrity, B=Benevolence, C=Competence$

α, β, γ denote the relative weights of the trust dimensions that sum up to 1.

The recommendation of an item to a user is computed as an aggregation of the recommendations of the n sources with the highest quality (weight), where the effect of each of the n sources of the final recommendation is relative to their weight (and therefore to the level of trust of the receivers). The recommendation function of an item i to a user u is defined as follows:

$$R(u, i) = \frac{\sum_{k=1}^n W(S_k, u) * r(S_k, i)}{n}$$

Where:

$R=recommendation, r=rating$

We expect that the relative importance of the trust dimensions will vary among different individuals based on their trust disposition. We also expect that the task domain (leisure Vs work related tasks) will also have an impact on the relative importance (weights) of the three dimensions of trust between receivers and sources. We plan to control for these variables in generating the recommendations. The model proposed above provides the first step towards the development of theoretical foundations for SF system design, so as to improve system performance. The next section describes how we plan to empirically validate the model proposed above.

4. Evaluation

4.1 Research Design

To evaluate the proposed model, we plan to conduct a 2x2 mixed design experiment. The first independent variable – the quality of the source – is manipulated at

two levels: i) control group where we would use the traditional SF proxy of source's trustworthiness – *Shared Interests* – as a baseline, and ii) treatment group where we include explicit ratings of the three trust dimensions - Integrity, Benevolence, and Competence – to generate recommendations. The second independent variable – task – is manipulated at two levels: i) leisure related task and ii) work related task.

We plan to recruit senior undergraduate students who know at least some other students reasonably well in both class rooms as well as university dormitories (to ensure the element of social relations). This would help us elicit trust beliefs about other participants in the study. Strict measures would be undertaken to ensure the confidentiality and privacy of the ratings (Not even the researcher would be able to identify the object of the ratings). For the first task we will generate a list of 50 popular movies, ensuring that these movies were watched by the subject population. For the second task we will generate homework questions and a set of relevant and non-relevant documents for these questions.

4.2 Procedures

The empirical study will be performed in two steps. First, we will bring together a group of subjects who are acquainted with one another, and ask each subject to choose 9 other subjects (i.e. sources)⁵ based on their ability to provide recommendations (3 who are expected to provide high-quality recommendations, 3 who are expected to provide low-quality recommendations, and 3 sources who the user is indifferent to their recommendations), and to assess the relationships to these sources on the three trust dimensions - *Integrity*, *Benevolence*, and *Competence*. In the second step, users will be

⁵ We restrict the number of people each subject has to assess to 9 only out of practicality; i.e. to limit the experiment duration.

asked to rate items which they have experienced. We will be using two different tasks from different domains to enable examination of the effect of the task. One task is related to entertainment. We will use a movie recommendation task on which users will be asked to rate movies, out of a list of popular movies. These ratings will serve (i) to calculate receiver-source *Shared Interests* (to be used in the baseline SF algorithm), and (ii) for making the recommendation. The other task is related to a professional aspect of the users. The users will have to rate documents relevant to a homework task in a university course. The system will then predict the relevancy of documents to the task based on the trust relations within the students. We will then examine whether these relations affect the quality of prediction. We will also be able to compare the effect of trust for different task domains, i.e. whether the tasks affect the model.

4.3 Operationalization

We intend to use the following operationalization for the proposed model's constructs.

- *Shared Interests*: similarity between the two vectors of user ratings, through standard measures used in SF (e.g. the Cosine function)
- *Trust Disposition*: McKnight's (2002) measure, adapted
- *Integrity*: McKnight's (2002) Integrity measure, adapted
- *Benevolence*: McKnight's (2002) Benevolence measure, adapted
- *Competence*: McKnight's (2002) Competence measure, adapted
- *Recommendation Quality*: based on standard measures of IF systems accuracy, (Herlocker et. Al., 2004)

4.4 System Implementation

We will split the data set into two sub-sets: the first will be used for training the model and calculating *Shared Interests*, while the second will be used to make predictions. We intend to test our model with a “golden standard” collaborative filtering system where the subjects’ ratings (of movies) are computed on a [-1, 1] scale (-1=strong negative; 1=strong positive). Each subject’s ratings will be normalized to account for difference between users

5. *Expected Contributions, Future Work, and Conclusion*

Our research adopted the design science (Hevner et al. 2004) approach to the development of filtering systems (i.e. by using theory to guide the design). The work-in-process described in this paper makes the first step towards this goal, namely in identifying the factors likely to impact recommendation quality and the design of an experiment to test their actual impact. We expect to be in a position to present the experiment’s findings by the conference date.

There are several limitations associated with the experimental design. First, the sample of subjects we intend to use may not represent fully the entire population of SF systems users. Second, we plan to use one “golden standard” collaborative filtering algorithm, while it is possible that alternative algorithms will produce different results. We plan to address these issues in the future as the current experiment is only a part of a larger project.

Recommendation systems, and specifically Social Filtering (SF) systems, play a significant role in reducing information overload and providing users with information relevant to their specific interest. For over a decade now, the ad-hoc standard in social

filtering employed an approach, where the shared-interests relation was used to generate recommendations, mainly due to difficulty in obtaining information about alternative types of receiver-source relations. The rapid growth in online social networks presents an opportunity for a new social filtering approach, where explicit relations from the social network could be utilized to indicate trusted sources. This approach could complement, or potentially even replace, the current practice of using similarity of interests to detect sources.

Bibliography

- Abdul-Rahman, A., Hailes, S., 1999, Relying on Trust to Find Reliable Information, In Proceedings 1999 International Symposium on Database, Web and Cooperative Systems (DWACOS'99), Baden-Baden, Germany.
- Benbasat I and Zmud R.W, The Identity Crisis within the IS Discipline: Defining and Communicating the Discipline's Core Properties, Management Information Systems Quarterly, June 2003, pp.183-194.
- Boger, Z., Kuflik, T., Shoal, P., and Shapira, B. "Automatic Keyword Identification by Artificial Neural Networks Compared to Manual Identification by Users of Filtering Systems", *Information Processing & Management*, 37 (2), 2001, pp. 187-198.
- Claypool , M.Brown, D., Le, P., Waseda, M. (2001) Inferring User Interest . In Proceedings of the 6th Intl. Conference on Intelligent User Interfaces (IUI 01), USA, pp. 33-40
- Freyne J. and Smyth B., 2004, An Experiment in Social Search, Proceedings of the 3rd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, Eindhoven, The Netherlands, pp. 95-103.

- Gnasa M., Won M., and Cremers A., 2004, Three Pillars for Congenial Web Searching – Continuous Evaluation for enhancing Web Search Effectiveness, *Journal of Web Engineering*, Vol. 3 No. 4, December, 2004, pp. 252-280
- Goldberg D., Nichols D., Oki B., and Terry D., 1992, Using Collaborative filtering to Weave an Information Tapestry, *Communications of the ACM*, Vol. 35 Issue 12, December 1992, pp. 61-70
- Guha R, Kumar R., Prabhakar R., and Tomkins A., 2004, Propagation of Trust and Distrust, in *Proceedings of the 13th Conference on world wide Web (WWW2004)*, May 2004, NY, USA
- Levin D., Cross R., and Abrams L., 2004, The Strength of Weak Ties You Can Trust: The Mediating Role of Trust in Effective Knowledge Transfer, *Management Science*, Vol. 40 No. 11, November 2004, pp. 1477-1490.
- Hanani, U., Shapira, B., and Shoval, P. "Information Filtering: Overview of Issues, Research and Systems", *User Modeling and User-Adapted Interaction-UMUAI*, 11, 2001, pp. 203-259
- Hardin R., 1993, The Street-Level Epistemology of Trust, *Politics and Society*, vol. 21, December 1993, pp. 505-529.
- Herlocker, J.L., Konstan, J.A., Terveen, L.G., and Riedl, J.R.T., 2004, "Evaluating Collaborative Filtering Recommender Systems", *ACM Transactions on Information Systems*, Vol. 22, No. 1, January 2004, pp. 5–53.
- Hevner, A., March, S., Park, J. and Ram, S., *Design Science in Information Systems Research, MIS Quarterly*, 2004, 28(1), pp. 75-105.
- Kelly, D., and Teevan, J. 2003. Implicit feedback for inferring user preferences, *SIGIR Forum* (37),2, pp. 18-28
- Kobsa A., 2001, Special Anniversary Issue, *User Modeling and User-Adapted Interactions*, 11(1-4).
- Mayer R., Davis J., and Schoorman F., 1995, An Integrative Model of Organizational Trust, *Academy of Management Review*, 20, pp. 709-734
- McKnight H., Choudhury V., and Kacmar C., 2002, Developing and Validating Trust Measures for e-Commerce: An Integrative Typology, *Information Systems Journal*, Vol. 13 Issue 3, pp.334-359.

- Morita, M., and Shinoda, Y. (1994). Information filtering based on user behavior analysis and best match text retrieval. In *Proceedings of the 17th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '94)*, Ireland, 272-281.
- Oard, D.W. 1997. The State of the Art in Text Filtering. *User Modeling and User-Adapted Interaction*, 7(3), pp. 141-178
- Shardanand U. and Maes P., 1995, Social Information Filtering: Algorithms for Automating “Word of Mouth, Proceedings of the SIGCHI conference on Human Factors in Computing systems, May 1995, Denver, Colorado, USA, pp. 210-217.