

# Social Filtering and Social Reality

Christopher Lueg

AI-Lab, Department of Computer Science  
University of Zurich

Winterthurerstrasse 190, CH-8057 Zurich

Email [lueg@ifi.unizh.ch](mailto:lueg@ifi.unizh.ch) Fax +41-1-63 56809

URL <http://www.ifi.unizh.ch/~lueg/>

## Abstract

In this paper, we argue that most current social information filtering approaches may benefit from more seriously taking into account the peculiarities of human cognition and human social behavior since current approaches only consider de-contextualized ratings. Social filtering systems exploit ratings provided by users in order to compute recommendations for other users. Typically, these ratings are detached from the situation and the social embedding in which they have been provided. Recent research on human cognition and behavior suggests that actions should not be viewed in isolation from the situation in which they occur (thus, the term “situated actions”). Accounting for the situation and the social embedding requires support for exploiting the situation rather than abstracting away the situation. In respect to exploiting the social embedding of ratings, we discuss the need for two related basic research directions. First, a self-organizing network of users trusting each other should be explored as a basis for true socially embedded filtering. Second, the suitability of collaborative filtering techniques as a tool for maintaining the focus of Usenet discussion groups by exposing spam and other clear off-topic postings should be investigated.

## Introduction

Social filtering systems, also referred to as collaborative filtering systems (Goldberg *et al.* 1992; Resnick *et al.* 1994; Konstan *et al.* 1997) or recommender systems (Resnick & Varian 1997), aim at automating the “word of mouth” (Shardanand & Maes 1995). Relying on recommendations given by others usually happens in situations with either too much or too few information available. Prime examples for successfully implemented social filtering processes are people reading newspapers since these people trust in the decisions of the editors to include the most interesting and important articles. Recommendations for movies, compact disks, books, and events given by editors of journals to help their customers or recommendations given by friends to help friends are other common examples

for relying on the judgments of others in unclear information situations.

Examples for the collaborative filtering approach, on the one hand, are systems filtering Usenet articles (e.g., Brewer & Johnson 1994; Resnick *et al.* 1994; Konstan *et al.* 1997; Terveen *et al.* 1997b). In the presence of a large amount of low quality items on the net, also called electronic junk (Denning 1982), the idea is that consumers help each other to distinguish between high quality and low quality items by providing ratings for items they have investigated. These ratings are collected and can then be used by others to focus on those items collectively rated best (or at least rated acceptable). Recommender systems, on the other hand, have been implemented in various domains, such as recommending webpages, music, or movies (e.g., Shardanand & Maes 1995; Terveen *et al.* 1997a). Despite less stressing the necessity of personal relations between the recommenders, the technique is basically the same as in the collaborative filtering approach.

Most social filtering approaches share some implicit assumptions that are explicated in the following. It seems as if the independence of ratings from both the topics and the representations of the objects being rated turns out to be the main lever applied by social filtering systems. Contrary to content-based filtering systems, social filtering systems are able to handle both virtual objects, such as Usenet articles or webpages, and real-world objects, such as movies or music, that are usually inaccessible to computers. In order to deal with ratings (a prerequisite for computing a recommendation) it is not necessary to analyze the corresponding objects as in content-based approaches. Also, the social embedding of recommendations can be abstracted away.

We proceed as follows: First, we briefly summarize why cognitive processes, such as rating, and socially embedded processes, such as recommending, cannot be replaced by “technical” processes without losing certain peculiarities. Next, we intro-

duce situatedness as a concept that appropriately accounts for the peculiarities of human cognition and briefly discuss “situated information filtering”. **spynews**, a newsreader that supports situated information filtering, supports situated actions by avoiding to abstract away the context in which the user’s “interest” occurs. This work has been focusing on individuals and their particular situation only. Social filtering as a community-based approach seems to be a promising complement to our individual-based filtering approach. Finally, in respect to appropriately accounting for the social embedding of ratings and recommendations, we discuss the need for two related basic research directions in social information filtering. First, a self-organizing network of users trusting each other may serve as a basis for true socially embedded filtering. Second, collaborative filtering techniques may be suitable for maintaining the focus of Usenet discussion groups by exposing spam and other off-topic postings.

### Socially Embedded Processes

Having computers imitating socially embedded processes, such as communication, collaboration, cooperation, negotiation, or recommendation, always raises a couple of important issues that have to be dealt with. Benefit is gained through automation since protocols and procedures can be handled more efficiently by computers compared to their human counterparts. However, if a social process is reduced to the exchange of tokens according to a protocol, the remaining process does not capture the *social* nature of the process involving mutual commitments, being under obligation, being responsible, etc. (Lueg & Müller 1996). The conceptualization of social processes as basically “technical” processes is in the tradition of the “rationalistic” perspective (Winograd & Flores 1986).

Put in a nutshell, the rationalistic perspective assumes that the world can be described objectively and that optimal (rational) solutions to problems can be deduced from these objective descriptions. Implications of the rationalistic perspective in the information filtering context are manifold (Lueg & Pfeifer 1997). For example, it is assumed that the “content” of a document can be observer-independently estimated on the basis of its representation. Also, it is assumed that “interest” can be estimated independently from the actual situation the recipient of information is involved in. Accordingly, it is assumed that ratings given by a particular person in a specific context can be appropriately represented in numeric ratings and that it makes sense to de-contextualize these ratings.

Regarding recommendations, the social context

of a recommendation is abstracted away from its social embedding; the recommendation is de-contextualized. Apparently, most current approaches to collaborative filtering are in the tradition of the rationalistic perspective.

### Situated Cognition

From a cognitive science and situated cognition perspective, the so-called “rationalistic” perspective does not appropriately capture human cognitive phenomena, such as cognition, knowledge, or behavior. Moreover, the rationalistic perspective does not provide an appropriate explanation for the notion of interest which is of outstanding relevance in the information filtering context. Contrary to the rationalistic perspective, which views human cognition as data-processing and behavior as being largely predetermined by plans, the situated cognition perspective suggests to view cognition, knowledge, and behavior as being fundamentally *situated*: cognition and knowledge are emergent properties of the interaction of an individual with its environment, i.e., its current situation (thus, the term “situatedness”). Cognition cannot be reduced to internal “data-processing”, it cannot be “de-contextualized” into a set of abstract descriptions (Suchman 1987; Clancey 1997). One important implication of situatedness is that the way a human interacts with a situation continuously changes based on his or her experience. Accordingly, we propose to view interest as being dynamically generated: interest is an emergent property of the interaction of an individual with an “information situation”.

Various approaches to find out about interest from different disciplines, such as psychology, information science, or computer science, can be found in the literature. Research on the notion of interest indicates that it is hard to determine why a specific document has actually been selected. Experiments (e.g., Lantz 1993; Mock 1996) have revealed that explanations of why a document was chosen for reading, or why it was found to be interesting varied and changed over time. The same result has been obtained when the subjects were asked about their initial information need. Situatedness explains why it is so hard to describe an information need. Information needs cannot be reduced to internal information processes alone, but require interaction with the current situation. Situational factors other than just the topical content of a selected document influence the relevance judgment. Factors influencing the judgment are any factors that the users bring into the situation, such as experience, background, knowledge level, beliefs, and personal preferences (Barry 1994). Also, the user’s judgment is influenced by the user’s purpose, the

user's expectation, the relevance of references, and future time savings (Su 1994). Accordingly, divergences between professional research judgments of relevance and precision, and actual user judgments have been reported in the literature (Su 1994).

### Situated Information Filtering

In general, the situated perspective applied to information filtering suggests that the goal is not to automate but rather to support information seeking processes in order to allow for situatedness and the peculiarities of human cognition. In an individual-based information filtering project, this perspective has led to the development of **spynews** (Lueg 1997), a newsreader supporting users in acting situated while browsing Usenet newsgroups. Instead of trying to find out about user interests as in traditional approaches, the newsreader monitors the user's newsreading behavior and uses a discussion-oriented approach to find out in what he or she is *not* interested. This allows **spynews** to filter uninteresting discussions in order to help the user focus on potentially interesting discussions.

It's a peculiarity of **spynews** that no model of interests is being constructed to draw inferences about the user's interests. Also, no content analysis of selected documents (Usenet articles in this context) is performed to find out why particular documents have been selected. **spynews** only reflects the user's behavior by gradually fading out uninteresting discussions. Since no model of interests is constructed, the situated information filtering approach avoids the abstraction problem that occurs when documents or user interests are formally described and compiled to profiles (Lueg 1998).

**spynews** has been implemented as an augmentation to the state-of-the-art **Knews**<sup>1</sup> newsreader. Preliminary tests with experienced Usenet users are encouraging. Additional extended user tests are under preparation in order to evaluate the benefits of this particular approach. So far, the **spynews** newsreader only tries to find out about in which discussions the user is *not* interested in. We extend the newsreader to provide additional hints to *interesting* discussions. In order to account for situatedness, these hints will also be based on the user's browsing behavior only. Examples for user actions that can reasonably be interpreted as indicators of interest are reading a particular discussion -partly or completely- or posting a followup article. Also, external user actions, such as sending email to a participant of a discussion or saving an article might be interpreted as indicators

---

<sup>1</sup><http://www.student.nada.kth.se/~su95-kjo/knews.html>

for interest in a discussion. However, *all* user actions are only *weak* indicators, since there are many other explanations that are equally plausible: The participation might be nothing more than a final statement and an (interesting) discussion might be ignored due to too much time pressure, or the user might want to think more about a topic before entering into the discussion, etc. (Lueg & Pfeifer 1997).

So far, our research on situated information filtering has been focusing on *individuals*. Applied to the *community-based* collaborative filtering approach, the situated perspective suggests that the social embedding of recommendations should be considered more seriously. A personal recommendation does not only depend on the particular situation of the recommender but also on the relation between the recommender and the recipient of the recommendation. Of course, editors of recommendations in journals hardly know their customers personally but they always have at least an idea of the target audience. It is yet unclear how this social embedding might be utilized in a general recommendation context. A practical example from the Usenet domain might help illustrate the social embedding of recommendations.

Discussions within Usenet on detecting "interesting discussions" showed that it is typically not only the topic of a discussion that influences whether the discussion is interesting or not. In addition, it is of outstanding relevance which persons contribute to a discussion. Although most people participating in the global, distributed conferencing system Usenet news do not know each other personally, one can observe a kind of emergent regard among the participants of a discussion group concerning the opinion of particular persons and the way they articulate their opinions. Interest in the people's opinions might even outvote a less interesting topic. The situated perspective suggests that exploiting this particular social embedding for filtering purposes requires a careful investigation of the issue. In what follows, we discuss several related issues.

### Future Research

Further research on information filtering and information overload situations is related to exploring the foundations for self-organizing "preference" networks, and investigating the usability of social filtering for spam-fighting and exposing clear off-topic postings. In the following, these issues are discussed in more detail.

**Networks of Trusted Users** Reports on experiences with Grouplens (Resnick *et al.* 1994), a collaborative filtering system for Usenet articles,

have shown that user acceptance is crucial especially at the beginning of a new collaborative filtering service since a critical mass of ratings is required for a working system (Miller, Riedl, & Konstan 1998). It has been argued that a kind of formal or implicit market system might be necessary to gain a sufficient number of ratings and to compensate those who consume ratings but do not provide ratings themselves (Konstan *et al.* 1997; Avery & Zeckhauser 1997). We investigate the development of tools supporting users in exchanging particular preferences with selected trusted people sharing interests.

Finding out about users sharing interests is a hot topic in collaborative filtering. The idea is that the preferences of one user with particular interests can be used as recommendations for other users with similar interests. However, if Usenet participants are viewed as situated agents that are embedded in a particular social environment (Usenet is best viewed as a virtual community), computing and comparing profiles in order to find out about shared interests turns out to be obsolete since people automatically find out about other people sharing their interests by participating in Usenet discussions. Trust in the judgments of others and regard to the opinions of others emerges the same way.

Familiarity with other Usenet participants is an emergent property of participating in Usenet discussions. This familiarity might be used as a basis for self-organizing networks of people trusting each other and exchanging profiles of likes and dislikes among them. We suspect that such a distributed network might provide sufficient social embedding to avoid the above mentioned motivational problems. Since Usenet itself is a self-organizing network of servers, chances are not too bad that such a preference network might be accepted within the Usenet community. Moreover, this distributed approach would avoid some of the resource problems that centralized approaches, such as the GroupLens system (Konstan *et al.* 1997), exhibit. Also, security problems would be less serious since interest profiles are only exchanged among users knowing and trusting each other.

### Exposing Spam and other off-topic postings

Net abuse is a hot topic within the global Usenet community. A collaborative filtering approach might turn out to be a powerful tool to fight spam and to expose off-topic postings. Spam<sup>2</sup> denotes the flooding of Usenet newsgroups with commercial advertisements. Negative effects of flooding newsgroups with spam (certain newsgroups exhibit up to ninety percent spam) are manifold. Users already using spammed newsgroups are driven away

since they increasingly have problems to detect new articles among uninteresting spam. While technically experienced participants may cope with spam by using sophisticated killfiles, new users not equipped with killfiles are kept away from spammed newsgroups.

Since hardly any participant in a newsgroup is interested in spam, keeping a newsgroup spam-free might provide enough motivation for the participants to provide ratings for a collaborative spam-filtering system. Besides having (seemingly) spam-free newsgroups, such a collaborative spam-fighting experiment would also provide valuable insights into the relation of varying interests among the participants of newsgroups and the motivational problems exhibited by traditional collaborative filtering systems. If such a collaborative spam-fighting system experiences a significantly better user acceptance than a traditional system, varying interests among the users of newsgroups could be identified as a reason for motivational problems keeping users from providing ratings in traditional collaborative filtering systems.

## Summary

Social filtering experiments in the Usenet domain have turned out to be less successful than expected. Motivational problems seem to prevent people from providing a sufficient number of ratings in order to bootstrap a successful collaborative filtering process. In this paper, we have argued that this failure may be due to not sufficiently considering situatedness and a lack of social embedding. Based on our work on a situated filtering approach focusing on individuals, we have pointed out various issues that should be treated more carefully in order to reach a higher degree of user acceptance. In addition, experiences with Usenet suggest that a self-organizing network of people exchanging preferences might be an alternative to centralized collaborative filtering approaches. Also, a collaborative filtering approach might turn out to be a powerful tool to fight net abuse, such as commercial advertisements flooding newsgroups (also called spam). In addition, collaborative filtering can help to maintain the focus of newsgroups by exposing off-topic postings.

## Acknowledgments

The authors would like to thank Michelle Hoyle and Ralf Salomon for their help, and Rolf Pfeifer for a stimulating research environment. We are grateful to Karl-Johan Johansson for the permission to use and to modify his **Knews** newsreader for our Usenet experiments.

---

<sup>2</sup><http://spam.ohww.norman.ok.us/default.htm>

## References

- Avery, C., and Zeckhauser, R. 1997. Recommender systems for evaluating computer messages. *Communications of the ACM* 40(3):88–89.
- Barry, C. 1994. User-defined relevance criteria: An exploratory study. *Journal of the American Society for Information Science (JASIS)* 45(3):149–159.
- Brewer, R., and Johnson, P. 1994. Collaborative classification and evaluation of Usenet. Technical Report CSDL-TR-93-13, Department of Information and Computer Sciences, University of Hawaii.
- Clancey, W. J. 1997. *Situated Cognition. On Human Knowledge and Computer Representations*. Cambridge University Press.
- Denning, P. J. 1982. Electronic junk. *Communications of the ACM* 25(3):163–165.
- Goldberg, D.; Nichols, D.; Oki, B. M.; and Terry, D. 1992. Using collaborative filtering to weave an information tapestry. *Communications of the ACM* 35(12):61–69.
- Konstan, J. A.; Miller, B. N.; Maltz, D.; Herlocker, J. L.; Gordon, L. R.; and Riedl, J. 1997. Applying collaborative filtering to Usenet news. *Communications of the ACM* 40(3):77–87.
- Lantz, A. 1993. How do experienced users of the system Usenet news select their information? Technical report, Department of Computer and Systems Science, University of Stockholm.
- Lueg, C., and Müller, M. 1996. Cooperative systems: The right direction? In *Proceedings of the Second International Conference on the Design of Cooperative Systems*, 315–329.
- Lueg, C., and Pfeifer, R. 1997. Cognition, situatedness, and situated design. In Marsh, J. P.; Nehaniv, C. L.; and Gorayska, B., eds., *Proceedings of the Second International Conference on Cognitive Technology*, 124–135. IEEE Computer Society.
- Lueg, C. 1997. An adaptive Usenet interface supporting situated actions. In *Proceedings of the 3rd ERCIM Workshop on User Interfaces for All*.
- Lueg, C. 1998. Supporting situated actions in high volume conversational data situations. In *Proceedings of the Annual ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'98)*. ACM Press.
- Miller, B. N.; Riedl, J. T.; and Konstan, J. A. 1998. Experiments with GroupLens: Making Usenet useful again. In *Proceedings of the 1997 USENIX Winter Technical Conference*.
- Mock, K. J. 1996. Hybrid hill-climbing and knowledge-based techniques for intelligent news filtering. In *Proceedings of the National Conference on Artificial Intelligence (AAAI'96)*. Menlo Park, California: AAAI Press.
- Resnick, P., and Varian, H. R. 1997. Recommender systems. *Communications of the ACM* 40(3):56–58.
- Resnick, P.; Iacovou, N.; Suchak, M.; Berstrom, P.; and Riedl, J. 1994. GroupLens: An open architecture for collaborative filtering of netnews. In Furuta, R., and Neuwirth, C., eds., *Proceedings of the International Conference on Computer Supported Cooperative Work (CSCW'94)*, 175–186. ACM Press.
- Shardanand, U., and Maes, P. 1995. Social information filtering: Algorithms for automating “word of the mouth”. In Irvin Katz et al., ed., *Proceedings of the Annual ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'95)*, 210–217. ACM Press.
- Su, L. 1994. The relevance of recall and precision in user evaluation. *Journal of the American Society for Information Science (JASIS)* 45(3):207–217.
- Suchman, L. 1987. *Plans and situated actions - The Problem of Human-Machine Communication*. Cambridge University Press.
- Terveen, L.; Hill, W.; Amento, B.; McDonald, D.; and Creter, J. 1997a. PHOAKS: A system for sharing recommendations. *Communications of the ACM* 40(3):59–62.
- Terveen, L.; Hill, W. C.; Amento, B.; McDonald, D.; and Creter, J. 1997b. Building task-specific interfaces to high volume conversational data. In *Proceedings of the Annual ACM SIGCHI Conference on Human Factors in Computing Systems (CHI'97) Conference Proceedings*, 226–233. ACM Press.
- Winograd, T., and Flores, F. 1986. *Understanding Computers and Cognition: A New Foundation for Design*. Ablex Publishing Corporation.