An Adaptive Usenet Interface Supporting Situated Actions

Christopher Lueg
AI-Lab, Department of Computer Science
University of Zurich
Winterthurerstrasse 190, CH-8057 Zurich
Tel. +41 1 63 54577 Fax +41 1 63 56809
lueg@ifi.unizh.ch

Abstract

In this paper, we discuss a novel approach to support users in information overload situations. This situated approach is based on the perspective that human cognition is most appropriately described as situated. In this context, situated means that human cognition is considered as an emergent property of the interaction of an individual with his or her environment. Our work aims at supporting humans in coping with information overload situations. In order to account for situatedness, we only support users in finding interesting information instead of automating the search task. Reading Usenet news is a prime example for situations with high volume conversational data. The global conferencing system Usenet offers an amount of conversational data per day that exceeds human cognitive capabilities by far. We have augmented a standard newsreading tool to offer support for situated actions. First experiences with the augmented newsreader are encouraging and suggest that this situated approach is an interesting complement to traditional information filtering approaches.

1 Introduction

The huge and increasing amount of information available in the information age suggests to investigate new ways to support humans in gathering information that might be interesting, helpful, or necessary for them. One of the central questions is how to provide adequate support for the information seeking process. From a cognitive science and situated cognition perspective, the goal is not to automate but to support this information seeking process in order to allow for situatedness and the peculiarities of human cognition.

In this paper, we present a situated information filtering approach to support users in coping with high volume conversational data. This approach is based on the perspective that human behavior is inherently situated (see below) and complements other approaches by avoiding to automatically detect information that might be interesting to the user.

The domain under investigation is the global conferencing system Usenet News offering more than 300000 articles per day. Support in coping with Usenet, on

the one hand, is given by a potentially significant[1] reduction of the amount of articles the user has to browse. On the other hand, gradually fading out of probably uninteresting discussions is used to support the user in becoming aware of his or her interests. Data collected while monitoring the user's information seeking behavior is used to find out about discussions that are likely to be *uninteresting* (instead of trying to determine the interesting discussions since this would involve modeling the user's interests). In order to avoid misinterpretations, the situated information filtering approach uses a high degree of interactivity. The augmented newsreader interface always allows the user to accept or to reject indicated consequences of his or her actions.

We will proceed as follows. First, we introduce the notion of situatedness and discuss some implications of this view for the design of human-computer interfaces and point out the major high-level design decisions. Based on these considerations, we discuss the benefits of the situated information filtering approach and describe how we modified a standard newsreader (`Knews`) to enable situated information filtering. The paper concludes with a discussion of first user experiences with the augmented newsreader.

2 Information Seeking and Situated Actions

From a situated cognition[4] perspective, user interests cannot not be viewed in isolation from human behavior since expressions of interest are only manifestations of human behavior. In contrast to the so-called "rationalistic" perspective[9] 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[14, 4]. One important implication of situatedness is that the way a human interacts with a situation continuously changes based on his or her experience.

Supposedly clear "expressions of interest", such as the selection of a specific document, are always subject to the frame-of-reference-problem[3]. The interest or information need is generated in the head of the observer, rather than in the head of the observed subject. The information need is not only "inside the user's head"[2], but a result of the interaction of the user with a continuously changing situation: the objectivity of a situation is achieved rather than given[14]. Accordingly, information-need situations are dynamic and constantly changing[1].

[1]The actual degree is dependent upon the particular newsgroup under consideration and the average length of its discussions.

3 Situated Information Filtering

Most approaches to help users in coping with high volume conversational data try to partly automate the information seeking process by matching models of interests and documents [6, 8, 11]. Since situated actions gain their meaning within the actual situation the user is involved (which is inaccessible to the newsreader), the situated information filtering approach does neither consider assumed *interests* of the users nor the *content* of selected documents.

Our approach focuses on supporting the user in acting situated rather than doing the information filtering for him or her. In our *discussion-oriented* approach, we help the user in focusing on more interesting discussions (also referred to as subjects or threads). Support is basically given by reducing the amount of *un*interesting discussions which are gradually faded out over time. As a sideeffect, the user is supported in becoming aware of his or her interests among the available discussions. In order to account for situatedness and the frame-of-reference problem, our approach completely avoids matching user profiles with document profiles (profiles in terms of content-based modeling of interests and documents).

Since Usenet consists of a dynamic stream of incoming articles, new articles have to be classified as belonging to more uninteresting or more interesting discussions. In our approach, we do not consider the *content* of the incoming articles but exploit the inherent structure of Usenet articles[5] to find out the discussions the articles belong to (see below).

In the literature, three different terms have been identified to describe filtering approaches: cognitive, social, and economic filtering[10]. The cognitive filtering approach characterizes the content of messages and the information needs of the recipient. The social filtering approach (also referred to as collaborative filtering[7] or recommending[13]) tries to exploit the interrelations of individuals in a community. The third approach, economic filtering, is only of minor interest in the context of information needs. In the content-based cognitive filtering approach, interests and document contents are abstracted to so-called profiles that are matched against each other to classify documents as supposedly interesting or uninteresting. In the social filtering approach, the contents of documents are not analyzed but expressions of interests in terms of abstract ratings are compiled to profiles of interests. In the use of abstraction, the social filtering approach is similar to the content-based one. Situated information filtering is different from these approaches since we avoid modeling in order to account for situatedness while the other approaches prefer the use of models.

The interaction with interfaces to information services, such as newsreaders for the Usenet or browsers for the World Wide Web, is technically highly constrained. Most interactions with the computer have to be performed with devices, such as a keyboard or a mouse, and most information is displayed on a screen only. Therefore, it is easy to monitor most of the user's interactions with the interface: the selection of a document, the delay between the selection and the selection of the next document (i.e., time *supposedly* spent to read documents), etc. These user actions are used to find out about in which discussions the user is *not* interested in.

Since the interpretation of situated (user) actions always involves a frame-of-reference issue[3], we have based the design of the augmented newsreader on the

following working hypotheses[9]:

1. Not too much value should be attributed to (single) user actions; they should not be interpreted as clear indicators of interest. (sometimes, humans make faults, or they are under the influence of events that do not directly relate to the document search; documents may turn out to be uninteresting after having selected them)

2. Selection of a document should not be interpreted as necessarily indicating an information need.

3. Not too much importance should be attributed to author, title, etc.

4. Information needs depend on the actual situation: they are dynamic, not static.

The global conferencing system Usenet News is a prime example for high volume conversational data since it offers more than 300.000 articles per day. Although the articles are already organized within a hierarchy of more than 15.000 so-called newsgroups (groups of articles sharing a particular topic, e.g., the newsgroup `comp.ai` has the topic computers and artificial intelligence), it is not uncommon that selected newsgroups still offer more than 1000 articles per day. Since most Usenet users participate in several newsgroups covering different topics, the overall article volume exceeds human cognitive capabilities (and the time people are willing to spend for using Usenet) by far.

For our Usenet experiments, we modified `Knews`[2], an X-windows-based newsreader that provides most standard newsreader features, such as threading, user-defined highlights and killfiles, and a graphical user interface. We modified the newsreader in order to monitor most user actions. Actions include, but are not limited to, selecting a newsgroup, selecting a thread or an article, reading an article partly or completely, saving or printing an article, forwarding or mailing an article, etc.

The detection of *un*interesting data can be accomplished by monitoring whether the user repeatedly ignores or even deletes a discussion while browsing a newsgroup. This ignoring of discussions is interpreted only as slight *indicator* for not being interested in the discussion. Therefore, ignoring a discussion does not cause the discussion to vanish immediately but to gradually fade out. Only if the discussion reaches its final "uninteresting" state (specifically visualized) it will be marked for exclusion for a user-defined period of time. Otherwise, in the next session, the discussion will be moved to the bottom of the list of unread discussions according to its current degree of "uninterestingness".

Different degrees of "uninterestingness" cause discussions behave like bubbles floating up and down according to the user's attention (in terms of actions) paid to the discussions: those with the most attention are listed at the top and those repeatedly ignored are located at the bottom and will eventually be filtered, i.e., they will not appear when the user enters the newsgroup next time. The filtering is based on the assumption that further articles (so-called followup articles) in uninteresting

[2] `http://www.student.nada.kth.se/~su95-kjo/knews.html`

discussions are also uninteresting[12]. However, things change and it is a peculiarity of Usenet discussions that subject changes often are not followed by a change of the corresponding `Subject:` line. Therefore, deletation of discussions is limited in time to allow for a reassessment of the degree of interestingness. This is necessary since the user is the only instance to ascribe interestingness in the situated filtering approach. In order to allow for the long-term discrimination of interesting and uninteresting discussions the behavioral state of discussions is collected in a database over extended periods of time.

4 Discussion

In this paper, we have presented a situated information filtering approach to support users in coping with high volume conversational data situations in the Usenet domain. The approach accounts for situatedness by avoiding the interpretation of user actions as clear indicators of interest. Instead, help is provided to focus the user's attention on potentially interesting articles by gradually fading out, deleting, and reordering discussions that have shown to be less interesting. Depending on the actual newsgroup, this approach significantly reduces the amount of uninteresting discussions and helps the user to focus his of her attention on the more interesting discussions.

Experiences with the prototypical situated filtering system have shown that it is indeed possible to provide an encouraging level of support without automating the information seeking process by matching models of user interests with documents. Experienced Usenet users find this more situated approach useful. However, further experiments have to be conducted to clarify whether this also holds true for less experienced or even novel users. Extended user experiments at our department are under preparation.

Although first results are encouraging, they should not be interpreted as if this purely situated approach is an *alternative* to traditional filtering and recommendation approaches. Instead, the situated approach should be regarded as a *complement* to traditional approaches. The golden mean is probably somewhere between a traditional filtering approach and this purely situated filtering approach.

Currently, our focus is on Usenet but we believe that the situated perspective is also valuable for related domains, such as email or the web that also exhibit "information overload". Also, the notion of situatedness introduces an important novel perspective into the technology-driven, largely rationalistic, Internet culture.

5 ACKNOWLEDGMENTS

The authors would like to thank Michelle Hoyle, Ralf Salomon, and Elke Siemon for their help and are grateful to Rolf Pfeifer for a stimulating research environment. Special thanks go to Karl-Johan Johnsson for the permission to use and to modify his `Knews` newsreader for our Usenet experiments.

References

[1] Carol L. Barry. User-defined relevance criteria: An exploratory study. *Journal of the American Society for Information Science (JASIS)*, 45(3):149–159, 1994.

[2] Nicholas J. Belkin and W. Bruce Croft. Information filtering and information retrieval: Two sides of the same coin? *Communications of the ACM*, 35(12):29–38, December 1992.

[3] William J. Clancey. The frame of reference problem in the design of intelligent machines. In K. van Lehn, editor, *Architectures for Intelligence. The 22nd Carnegie Mellon Symposium on Cognition*, pages 257–423. 1991.

[4] William J. Clancey. *Situated Cognition. On Human Knowledge and Computer Representations.* Cambridge University Press, in press.

[5] M. Horton and R. Adams. RFC-1036 (request for comments), December 1987.

[6] Andrew Jennings and Hideyuki Higuchi. A personal news service based on a user model neural network. *IE-ICE Transactions on Information Systems*, March 1992.

[7] Joseph A. Konstan, Bradley N. Miller, David Maltz, Jonathan L. Herlocker, Lee R. Gordon, and John Riedl. Applying collaborative filtering to usenet news. *Communications of the ACM*, 40(3):77–87, March 1997.

[8] Ken Lang. Newsweeder: Learning to filter netnews. In *Proceedings of the Twelth International Conference on Machine Learning.* Morgan Kaufmann, 1995.

[9] Christopher Lueg and Rolf Pfeifer. Cognition, situatedness, and situated design. In Jonathan P. Marsh, Chrystopher L. Nehaniv, and Barbara Gorayska, editors, *Second International Conference on Cognitive Technology*, pages 124–135. IEEE Computer Society, 1997.

[10] Thomas W. Malone, Kenneth R. Grant, Franklyn A. Turbak, Stephen A. Brobst, and Michael D. Cohen. Intelligent information-sharing systems. *Communications of the ACM*, 30(5):390–402, May 1987.

[11] Kenrick J. Mock. Hybrid hill-climbing and knowledge-based techniques for intelligent news filtering. In *National Conference on Artificial Intelligence*, Menlo Park, California, 1996. AAAI Press.

[12] Masahiro Morita and Yoichi Shinoda. Information filtering based on user behavior analysis and best match text retrieval. In *International ACM-Conference on Research and Development in Information Retrieval (SIGIR 94).* Springer-Verlag London, 1994.

[13] Paul Resnick and Hal R. Varian. Recommender systems. *Communications of the ACM*, 40(3):56–58, March 1997.

[14] Lucy Suchman. *Plans and situated action - The Problem of Human-Machine Communication.* Cambridge University Press, 1987.