

Implicit Feedback for Recommender Systems

From: AAAI Technical Report WS-98-08. Compilation copyright © 1998, AAAI (www.aaai.org). All rights reserved.

Douglas W. Oard and Jinmook Kim

Digital Library Research Group
College of Library and Information Services
University of Maryland, College Park, MD 20742
{oard, jinmook}@glue.umd.edu

Abstract

Can implicit feedback substitute for explicit ratings in recommender systems? If so, we could avoid the difficulties associated with gathering explicit ratings from users. How, then, can we capture useful information unobtrusively, and how might we use that information to make recommendations? In this paper we identify three types of implicit feedback and suggest two strategies for using implicit feedback to make recommendations.

Introduction

Recommender systems exploit ratings provided by an entire user population to reshape an information space for the benefit of one or more individuals (Oard, 1997b). In research systems, these ratings are often provided explicitly by each user using one or more ordinal or qualitative scales. The cognitive load effort to assign accurate ratings acts as disincentive, making it difficult to assemble large user populations and contributing to data sparsity within existing populations. Implicit feedback techniques seek to avoid this bottleneck by inferring something similar to the ratings that a user would assign from observations that are available to the system. Such an approach could greatly extend the range of applications for which recommender systems would be useful.

Sources of Implicit Feedback

Nichols (1997) surveyed the state of the art in implicit feedback techniques with an eye toward their potential use for information filtering. *Table 1* presents the sources identified by Nichols and some others that we believe will also be useful.¹ In addition to explicit ratings we have identified three broad categories of potentially useful observations: examination, retention and reference.

Information systems often provide brief summaries of several promising documents using some sort of selection interface display, and selection of individual objects for further examination can thus provide the first cue about a

¹ Nichols (1997) suggested two additional behaviors related to content-based retrieval: discovery of users that present a common set of query terms and discovery of users that retrieve similar documents. Both can be mapped into our framework by adopting the perspective that queries are information objects in their own right.

user's interests. USENET newsreader software typically records the identifiers of messages that users have seen, and Karlgren (1994) explored the design of a recommender system using such lists. Morita and Shinoda (1994) and Konstan et al. (1997) found a positive correlation between reading time and explicit ratings in USENET news applications, and we have generalized that source of observations as "examination duration" to accommodate other modalities such as audio and video. Hill et al. (1992) have developed this idea further, defining "edit wear" as an analogue to the useful effects of uneven wear that physical materials accumulate over time that provide other users with cues that help discover useful materials and useful items of those items. In text browsing, for example, edit wear might be measured by using dwell times at specific locations in the text to characterize scrolling behavior. Examination may extend beyond more than a single interaction between user and system, and we seek to capture that source of observations by characterizing the repetition of the foregoing user behaviors. Finally, when information access is priced on a per-item basis, purchase decisions offer extremely strong evidence of the value ascribed to an object. Similar information would be available at a somewhat coarser scale when users purchase subscription access to certain types of content (e.g., subscription to a separately priced cable television channel).

Category	Observable Behavior
Examination	Selection Duration Edit wear Repetition Purchase (object or subscription)
Retention	Save a reference or save an object (with or without annotation) (with or without organization) Print Delete
Reference	Object->Object (forward, reply, post follow up) Portion->Object (hypertext link, citation) Object->Portion (cut & paste, quotation)

Table 1. Observable behavior for implicit feedback

Our “retention” category is intended to group those behaviors that suggest some degree of intention to make future use of an object. Bookmarking a web page is a simple example of such a behavior, and we have generalized that idea as “save a reference” to accommodate a wider range of actions such as construction of symbolic links within a file system. Rucker & Polanco (1997), for example, constructed a recommender system using bookmark lists. Saving the object itself is the obvious alternative, something Stevens (1993) used as implicit feedback for content-based filtering. In either case, the object may be saved with or without some form of annotation. For example, web browsers typically default to using the page title in the bookmark list, but users may optionally provide a more meaningful entry if they desire. Although numerous confounding factors would likely be present, it may be possible to infer something about the value a user places on an individual page by whether or not they go to the trouble of constructing an informative bookmark entry. Similarly, users may choose to save a reference or an object in an explicitly organized fashion or in the default manner. For example, storing electronic mail about this workshop in a new folder might provide greater support for an inference that the user ascribes particular value to the message than would the use of some default scheme such as placing it in the folder routinely used for mail from the message’s originator. The salient issue in this case is not the act of organizing, but rather the way in which the organization given to an individual object distinguishes it from the way in which similar forms of organization are assigned to other objects. This difference may not be easy to characterize, but it may be worth thinking about how to do it. We have chosen to group printing with retention because of the permanence of the printed page, but users may also print document or images to facilitate examination because paper still has some decided advantages over electronic displays in many applications. Printing overlaps with the next category (reference) as well, since users may print a document or image with the intention of forwarding them to another individual or including portions in another document. Nevertheless, printing is often associated with a desire for retention, so we find this grouping useful. As with examination, it may be possible to infer something about the portions of a document that the user finds most valuable from the portions which he or she chooses to print. Finally, the retention category is distinguished by the possibility of directly observing evidence of negative evaluations as well. When retention is a default condition, as in some electronic mail systems, a decision by the user to delete an object might support to an inference that the deleted object is less valued than other objects that are retained.

The “refer to” category may appear at first glance to contain a fairly eclectic group of observable activities, but each has the effect of establishing some form of link between two objects. Forwarding a message, for example, establishes a link between the new message and the original. Similarly, replying individually or posting a follow up

message to some form of group venue such as a mailing list establishes the same sort of link. Goldberg et al. (1987) described a simple example of this in which users could construct an electronic mail filter to display messages that their colleagues had taken the time to reply to. Hypertext links from one web page to another and bibliographic citations in academic papers create links from a portion of an object (characterized, perhaps, by some neighborhood around the link itself) to another object, although the refinement to a portion of a document has not been exploited often. Brin & Page (1998) provide an example of how hypertext links might be used, although their focus is on a population statistics rather than individual preferences. Garfield (1979) describes the design of retrieval systems that are based on bibliographic citations. Alternatively, selective inclusion of another document, using either cut-and-paste or a quotation, creates a link from an information object to a portion of another.

Using Implicit Feedback

The goal of a recommender system is to help users find desirable information objects. That task combines inference and prediction, and Figures 1 and 2 show alternative strategies for accomplishing this. Figure 1 depicts a modular strategy in which the inference stage seeks to produce ratings similar to those that a user would have explicitly assigned, and then the prediction stage uses those estimated ratings to predict future ratings. Konstan et al. adopted this perspective when evaluating how well observed reading time predicted explicit ratings for individual articles. Figure 2 shows an alternative strategy in which past observations are used to predict user behavior in response to new information, and then the inference stage seeks to estimate the value of the information based on the predicted behavior. We are not aware of any implementations of this second approach, but Stevens (1993) implemented a simplified version of the strategy. He predicted the examination duration for a new USENET news article based on the examination durations for similar articles in the past and then constructed content-based queries that would select articles with long predicted examination durations. This essentially amounts to a degenerate inference stage in which desirability is assumed to increase monotonically with examination duration.

The distinction between the two strategies is quite subtle in the case of content-based filtering. In a recommender system, by contrast, the strategy shown in Figure 1 would characterize each article using the examination durations reported by other users, while the strategy shown in Figure 2 would characterize each article using the predicted ratings for other users. Recommender systems based on the second strategy might be more flexible, since participating users might draw different inferences from the same observations if they did not share a common set of objectives. On the other hand, recommender systems using the first strategy would likely have more context available locally for interpreting observations than would be available at

other points in the network. It might thus be worth considering hybrid approaches in which some preliminary interpretation is performed locally when the observation is made and then additional inferences are drawn at other points in the network.

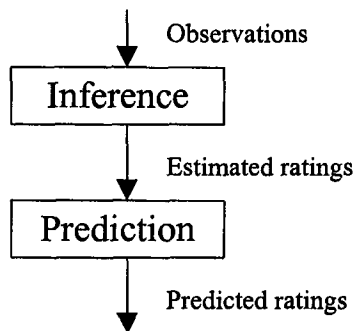


Figure 1. Rating estimation strategy.




Figure 2. Predicted observations strategy.

Conclusion

We have presented three potential sources for implicit feedback and described two ways those sources could be used by recommender systems. Our “examination” category seeks to capture ephemeral interactions that begin and end during a single session, while the “retention” category groups user behaviors that suggest an intention for future use of the material. Our third category is reference, which includes user behaviors that create explicit or explicit links between information objects. We believe these categories group observable behavior in a way that is useful when thinking about how to make predictions, and toward that end we have suggested two strategies for using implicit feedback in recommender systems. Our present work is focused on understanding how to relate observations to predicted ratings, both individually and in various combinations that could be more informative than single-source observations. We then hope to develop and implement a prototype that will give us some insight into how implicit feedback can be used effectively in an application environment. If successful, this approach could help transcend the current reliance on explicit ratings and thus signifi-

cantly expand impact and importance of recommender systems in a networked world.

References

- Brin, S. and Page, L. 1998. The Anatomy of a Large-Scale Hypertextual Web Search Engine. Dept. of Computer Science, Stanford Univ. <http://google.stanford.edu/~backrub/google.html>.
- Goldberg, D., Nichols, D., Oki, B. M., and Terry, D. 1992. Using Collaborative Filtering to Weave an Information Tapestry. *Communication of the ACM*, December, 35(12): 61-70.
- Garfield, E. 1979. *Citation Indexing: Its Theory and Application in Science, Technology, and Humanities*. New York: Wiley-Interscience.
- Hill, W.C., Hollan, J. D., Wrobelwski, D. and McCandless, T. 1992. Read Wear and Edit Wear. In: *Proceedings of ACM Conference on Human Factors in Computing Systems, CHI '92*: 3-9.
- Karlgren, J. 1994. *Newsgroup Clustering Based on User Behavior: A Recommendation Algebra*. Technical and Research Reports from SICS, T94-01. <http://www.sics.se/libabstracts.html#T94/04>.
- Konstan, J. A., Miller, B. N., Maltz, D., Herlocker, J. L., Gordon, L. R., and Riedl, J. 1997. GroupLens: Applying Collaborative Filtering to Usenet News. *Communications of the ACM*, March, 40(3): 77-87.
- Morita, M. and Shinoda, Y. 1994. Information Filtering Based on User Behavior Analysis and Best Match Text Retrieval. In *Proceedings of the Seventeenth Annual International ACM-SIGIR Conference on Research and Development in Information Retrieval*, 272-281.
- Nichols, D. M. 1997. *Implicit Ratings and Riltering*. In *Proceedings of the 5th DELOS Workshop on Filtering and Collaborative Filtering*, 10-12. Budapaest, Hungary, ERCIM.
- Oard, D. W. 1997. The State of the Art in Text Filtering. *User Modeling and User-Adapted Interaction*, 7(3): 141-178. <http://www.glue.umd.edu/~oard/research.html>.
- Rucker, J. and Polanco, M. J. 1997. Personalized Navigation for the Web. *Communications of the ACM*, March, 40(3): 73-89.
- Stevens, C. 1993. *Knowledge-Based Assistance for Accessing Large, Poorly Structured Information Spaces*. Ph.D. dissertation, Dept. of Computer Science, Univ. of Colorado, Boulder. <http://www.holodeck.com/curt/mypapers.html>.