

Learning about Users from Observation

Ingo Schwab, Alfred Kobsa, Ivan Koychev

GMD FIT.MMK
D-53754 Sankt Augustin, Germany

{Ingo.Schwab,Alfred.Kobsa,Ivan.Koychev}@gmd.de

Abstract

Many approaches and systems for recommending information, goods, or other kinds of objects have been developed in recent years. In these systems, machine learning methods are often used that need training input to acquire a user interest profile. Such methods typically need positive and negative evidence of the user's interests. To obtain both kinds of evidence, many systems make users rate relevant objects explicitly. Others merely observe the user's behavior, which yields positive evidence only; in order to be able to apply the standard learning methods, these systems mostly use heuristics to also find negative evidence in observed behavior.

In this paper, we present an approach for learning interest profiles from positive evidence only, as it is contained in observed user behavior. Thus, both the problem of interrupting the user for ratings and the problem of somewhat artificially determining negative evidence are avoided.

A methodology for learning explicit user profiles and recommending interesting objects has been developed. It is used in the context of ELFI – a Web-based information system. The evaluation results are briefly described in this paper.

Our current efforts revolve around further improvements of the methodology and its implementation for recommending interesting web pages to users of a web browser.

Introduction

In the last few years, many approaches and systems for recommending information, products and other items have been developed. These systems try to help users find pieces of information or other objects in which the users will presumably be interested [3]. Since the Internet is growing extremely large with a vast amount of information accessible to anyone with a computer, recommender systems are becoming increasingly more important.

Two main different approaches for recommending objects to the user have been developed so far. Feature-based filtering systems take individual preferences with respect to certain features of objects into account (in "content-based" information filtering, the content is described by a restricted number of characteristic features of the content,

e.g. characteristic words). Clique-based (aka "collaborative") filtering systems instead typically build on similarities between users with respect to the objects in which users implicitly or explicitly express an interest.

Machine learning methods can be used to solve classification problems. Hence, a straightforward way of using machine learning for acquiring interest profiles is to assume that the set of information objects can be divided into classes (e.g., for "interesting" and "not interesting"). In many systems, users must provide examples for both classes in an initial training phase, on the basis of which a classification algorithm is learned inductively. Thereafter, the classification algorithm can determine whether new information objects belong to the "interesting" or to the "not interesting" class. Such explicit rating requires additional user effort and keeps users from performing their actual tasks, both of which is undesirable. As has been observed by Carroll and Rosson [1], users are unlikely to engage in such additional efforts even when they know that they would profit in the long run. Additionally, motivating Web consumers to provide personal data is proving very difficult. Internet users normally avoid engaging in a relationship with Internet sites. This is mostly due to a lack of faith in the privacy policy of today's web sites. Users either withhold personal data or provide false data¹². Conclusions about user interest should therefore not rely very much on user ratings, but rather take passive observations about users into account as far as possible.

¹<http://www.thestandard.com/article/display/0,1151,235,00.html>

²http://www.cc.gatech.edu/gvu/user_surveys/survey-1998-10/graphs/privacy/q48.htm

Our goal is to develop a content-based recommendation component. In order to be unobtrusive, it shall learn individual interest profiles based on passive observation only. The central source of information about users' interest is their web navigation, i.e. the sequence of users' movements on the Web. In each navigation step, users select among the currently selectable objects. In similar situations, other systems use heuristics to determine positive and negative evidence of the users' information interest (e.g., unselected objects are counted as negative examples [4][6]). For a general approach, however, we claim that unselected objects which are interesting to the user may exist (they may have just been overlooked or will perhaps be visited later). Classifying them as negative examples is dangerous since many of these classifications may be wrong and cause too much noise in the training data. It is more suitable to only take selected objects as examples for the "interesting" class, and to disregard objects that have not been selected. Unfortunately, though, in this case standard classification methods are not applicable. For learning interest profiles, we therefore had to invent new learning methods or revise existing ones.

Related Work

In the past, several systems have been developed that employ learning procedures to identify individual users' interests with respect to information objects and their contents, and make use of this interest profile to generate personalized recommendations.

Lieberman [4] developed the system Letizia, which assists a user in Web browsing. It tries to anticipate interesting items on the Web that are related to the user's current navigation context (i.e., the current Web page, a search query, etc.). For a set of links it computes a preference ordering, based on a user profile. This profile is a list of weighted keywords, which is obtained by aggregating the results of TFIDF analyses of pages. Letizia uses heuristics to determine positive and negative evidence of the user's information interest. Viewing a page indicates interest in that page, bookmarking a page indicates even stronger interest, while "passing over" links (i.e., selecting a link below and/or on the right of other links) indicates disinterest in these links.

A classification approach is taken by Syskill&Webert [7]. The user rates a number of

Web documents from some content domain on a binary "hot" and "cold" scale. Thus, positive and negative learning examples become available to the system. Based on these ratings, it computes the probabilities of words being in hot or cold documents. A set of word probability triplets is formed for each user, which can be regarded as an interest profile that characterizes the average hot and cold documents of this user. Based on this profile, the Naive Bayes Classifier method is used to classify further documents as hot or cold, respectively.

The system Personalized WebWatcher [6] also uses the Naive Bayes Classifier. This system watches individual users' choices of links on Web pages, in order to recommend links on other Web pages that are visited later. The user is not required to provide explicit ratings. Instead, visited links are taken as positive examples, non-visited links as negative ones.

The Naive Bayes Classifier is again used in the system NewsDude [2], similarly to Syskill&Webert, to recommend news articles to users. In NewsDude, the probabilities are taken to characterize the long-term interests of a user. To avoid recommending too many similar documents to a user, an additional short-term profile is built by memorizing currently read articles. New articles are then compared to the memorized ones; if they are too similar, they are not recommended although they will typically match the long-term interest profile. This procedure corresponds to the nearest-neighbor classification algorithm, which is well known in Machine Learning. Note that for the short-term profile, positive examples are only needed (albeit to produce "negative" recommendations).

Learning about the User

A standard approach when using machine learning for acquiring interest profiles is to assume that the set of information objects can be divided into classes (e.g., "interesting" and "not interesting"). Then, content-based filtering can be performed by letting the user classify given documents, thus providing examples for both classes (see e.g. [6]), and apply an inductive classification algorithm to these examples. For further documents, the classification algorithm can determine whether they belong to the "interesting" or to the "not interesting" class.




Figure 1 Prediction accuracy of compared approaches according to the number of selected documents.

In general, it can be assumed that it is legitimate to divide the objects from the information subspace into two such classes. People are normally interested or disinterested in given topics. However, obtaining an appropriate set of negative examples (i.e., examples of the “not interesting” class) is problematic. The central source of information about the user is his or her web navigation behavior, and especially the set of selected objects. Selections are made from the set of currently available web pages. We already mentioned that there are systems which use unselected objects as negative examples. However, for a general approach we claim that unselected objects may exist that are nevertheless interesting to the user. Since there are millions of pages on the Internet, it is a common situation that pages are overlooked, and it is impossible to have an overview over all relevant pages. Sometimes pages that are not visited at the moment become visited at a later point, and sometimes they are ignored forever even when the user is interested in them since it is too time consuming or simply not possible to follow every interesting link. Therefore, classifying the objects not visited as negative examples seems to be a dangerous assumption. It may happen that wrong decision borders are calculated between the classification regions. It is more suitable to only

take selected objects as examples for the “interesting” class. However, in this case standard classification methods are not applicable. Thus, for learning interest profiles we had to invent new learning methods or revise existing ones. This section can only give a brief survey about the different methods studied (for a more detailed discussion see [8]).

In our project, a probabilistic approach and an instance-based learning approach [5] have been used. They can be applied to learn a general characterization of objects that are relevant to the user. We modified both approaches to deal with a single class only by employing a notion of similarity or distance between objects. However, it is difficult to use these learning results to characterize individual users’ preferences explicitly, which is a desirable feature of user modeling systems [3]. Therefore, we developed a third mechanism that aims at selecting those features that are extraordinarily important to the user for identifying relevant objects.

It turned out that this feature selection method additionally helps improve the distance measure for instance-based learning (see Fig. 1). Moreover, feature selection can be combined with both probabilistic and instance-based learning to focus

the learning task. The developed algorithms have been implemented and evaluated [8] in the real-world application ELFI (ELeCtronic Funding Information). ELFI¹ is a WWW-based information system that provides information about research grants and is used by more than 1000 users in German research organizations who monitor and/or advise on extra-mural funding opportunities. In this system, additional calls for proposals are recommended based on those that the user had already browsed so far.

The conclusions of our experiments are twofold. First, our experiments demonstrate that the use of feature selection significantly improves the performance of the learning algorithms. Instance-based learning plus feature selection works well for small training sets even with a simple Hamming distance. However, with growing training set it becomes apparent that weighted distance measure learns much faster. Second, instance-based learning performs better than the probabilistic approach.

Since we are not only interested in a single adaptivity task (i.e., predicting user-specific degrees of object relevance) but also in determining explicit information about user interest and/or preferences, we employed statistical methods to find the object features that are especially important to an individual user. While instance-based and probabilistic methods try to assess objects as a whole when determining their interestingness, this latter approach results in interest degrees for selected features that characterize the user (instead of the objects). This is more in line with traditional user modeling approaches where user models are knowledge bases with explicit representation of user characteristics.

Current Research

Our goal was to develop methods that are able to passively observe users and, based on positive examples only, recommend objects which are presumably also interesting for the user. The methods have been implemented in the ELFI environment and the empirical results were very encouraging.

Our current research interests lie in the improvement of the developed methodology and in its application to web browsing. It shall observe

user actions in the WWW. The trace of user selections, scrolling and navigation operations in a web browser, the Web Navigation, is used as a source of implicit information about user interests. Information retrieval techniques (e.g. TFIDF) and feature selection will analyze the visited pages. After that, an explicit user profile representing user interests and disinterest is generated. Finally, the system searches for pages on the WWW fitting the user's interests. Modified machine learning algorithms for learning from positive examples only will be used for selecting and recommending pages that are the most relevant to the current user's interests.

Even though it is a different application, the main problem remains the same: to extract users' interest from the content of the visited web pages to recommend other pages that are relevant to her current interests. We anticipate that the developed algorithms should also work in this new environment.

Conclusion

We have shown that it is possible to learn user interests and disinterest implicitly. Though we do not claim that this approach produces better results than explicitly asking the user about her interests, in many applications it is often not possible to receive user ratings. In some cases it is too time consuming or simply impossible for the user to give feedback. In other cases the user does not trust the system and is unwilling to reveal his or her personal interests. And as a general observation, users are extremely reluctant to perform any actions that are not directed towards their immediate goal (like training the learning algorithm) if they do not receive immediate benefits.

For developers of intelligent interactive systems, it remains a challenge to design interfaces that can acquire user feedback in an unobtrusive way so that implicit user ratings will become more easily available (some special cases where this seems possible are discussed in [3]). However, we think that in cases where users have to select interesting objects from larger sets, negative evidence will always be difficult to obtain. In these cases, the methods presented in this paper can be fruitfully employed.

¹ <http://www.elfi.ruhr-uni-bochum.de/elfi/>

References

- [1] Carroll, J. and Rosson, M. B. 1987. The paradox of the active user. In J. M. Carroll (Ed.), *Interfacing Thought: Cognitive Aspects of Human-Computer Interaction*. Cambridge, MA: MIT Press.
- [2] Billsus, D., and Pazzani, M. J. 1999. A Hybrid User Model for News Classification. In *Kay J. (ed.), UM99 User Modeling - Proceedings of the Seventh International Conference*, pp. 99-108. Springer-Verlag, Wien, New York.
- [3] Kobsa, A., Koenemann, J. and Pohl, W. (2000). *Personalized Hypermedia Presentation Techniques for Improving Online Customer Relationships*. Forthcoming.
- [4] Lieberman, H. 1995. Letizia: An Agent That Assists Web Browsing. *International Joint Conference on Artificial Intelligence*, Montréal.
- [5] Mitchell T. Instance-Bases Learning. Chapter 8 of *Machine Learning*. McGraw-Hill, 1997.
- [6] Mladenic, D. 1996. Personal WebWatcher: Implementation and Design. Technical Report IJS-DP-7472, Department of Intelligent Systems, J. Stefan Institute, Slovenia.
- [7] Pazzani, M. J. and Billsus, D. 1997. Learning and Revising User Profiles: The Identification of Interesting Web Sites. *Machine Learning*, 27, 313-331
- [8] Schwab, I., Pohl, W., and Koychev, I. 2000. *Learning to Recommend from Positive Evidence*, Proceedings of Intelligent User Interfaces 2000, ACM Press, forthcoming.