Greedy Recommending Is Not Always Optimal

Maarten van Someren[1], Vera Hollink[1], and Stephan ten Hagen[2]

[1] Dept. of Social Science Informatics, University of Amsterdam,
Roetersstraat 15, 1018 WB Amsterdam, The Netherlands,
{maarten,vhollink}@swi.psy.uva.nl
[2] Faculty of Science, University of Amsterdam,
Kruislaan 403, 1098 SJ Amsterdam, The Netherlands
stephanh@science.uva,nl

Abstract. Recommender systems suggest objects to users. One form recommends documents or other objects to users searching information on a web site. A recommender system can use data about a user to recommend information, for example web pages. Current methods for recommending are aimed at optimising single recommendations. However, usually a series of interactions is needed to find the desired information. Here we argue that in interactive recommending a series of normal, 'greedy', recommendings is not the strategy that minimises the number of steps in the search. Greedy sequential recommending conflicts with the need to explore the entire space of user preferences and may lead to recommending series that require more steps (mouse clicks) from the user than necessary. We illustrate this with an example, analyse when this is so and outline when greedy recommending is not the most efficient.

1 Introduction

Recommender systems typically recommend one or more objects that appear to be the most interesting for a user. A large number of methods have been proposed and a number of systems have been presented in the literature, e.g. [1, 2,8,4,5,6,9,10,16]. These systems collect information about the user, for example documents, screens or actions that were selected by the user, and use this to recommend objects. Some authors view 'return on investment' as the criterion for success of recommendings (e.g. [7]). In this case recommending is a form of advertising and the economic value of objects is of key importance because this will determine the 'return on investment'. A related but different criterion is 'user satisfaction with the recommendation and the recommended object'. A user may be most satisfied with an object that has little 'return on investment' for the site owner but that has a high value for the user. In this case the goal of the recommender can be to maximise either 'return on investment' or user satisfaction. Another dimension of user satisfaction is effort in using the site, for example the number of clicks. A recommender can have as its goal to minimise user effort, for example in a situation where the user will eventually find his target object or information. Of course a recommender can also aim to maximise some

B. Berendt et al. (Eds.): EWMF 2003, LNAI 3209, pp. 148–163, 2004.

- Goal
  - Benefit site owner
  - Benefit user
- Criterion
  - Maximise value of object
  - Minimise effort of finding it
- Preferences distribution (per user)
  - One target and rest flat
  - One target and rest (partially) ordered
  - Multiple targets

Fig. 1. Dimensions of recommender systems

combination of user effort and value of the result. These different recommending tasks require different methods. Advertising can be viewed as a kind of game in which the user and the vendor pursue their own goals and in maximising user satisfaction, user and site owner share the same goal.

If user satisfaction is the goal then another dimension of the recommending task is important: is the users goal a single object or are there many objects that will satisfy the user, as for a user who is just surfing. In the second case, the main goal of recommending is to suggest useful objects, for example something unexpected. If the user has a specific goal then recommending is similar to information retrieval. The purpose of recommending in this case is to help the user to find an object that maximally satisfies the users goal and to minimise his effort in finding it. Information retrieval is normally based on user-defined queries but in some applications users are not able to formulate adequate queries because they are not familiar with the domain, the terminology and the distribution of objects. In this case presenting specific objects can replace or complement a dialogue based on queries. Figure 1 summarises the main dimensions of recommending tasks.

If the criterion is to minimise effort then the choice of an object to recommend depends on two different goals: (1) to offer candidate objects that may satisfy the users interest and (2) to obtain information about the users preferences. A single object may be optimal for both goals but this is not necessarily the case. In this paper we show that different recommendations may be optimal for these goals and that recommending the best candidate ('greedy recommending') does not always minimise the number of user actions before the target is found. We demonstrate this by introducing an alternative method that exploits a particular type of pattern in user preferences, requiring on average fewer user actions than greedy recommending.

By analogy with greedy heuristic search methods, we use the term *greedy* recommending when the recommender presents objects that it predicts to be closest to the target. This can be seen as a myopic decision making process in which the recommender aims at offering the user immediately the object of interest. If recommending takes information about the user (like the interaction history)