



Web Dynamics (SS 10) Assignment 7

Handout on: July 1, 2010

Due on: July 8, 2010

Exercise 7.1: Similarity of users in social networks

The lecture considered only simple similarity measures for users based on their votes on items. In real systems, a lot more information is available. Consider a social network where users collect sets of items, assign tags to them, and maintain a network of friends. Canonical examples for such social networks are del.icio.us (where the items are bookmarks), librarything (books), flickr (images), or bibsonomy (scientific articles).

Discuss extensions of the simple user similarity models for such networks. Think of integrating items, content of items, tags, and friends into the similarity measure.

Exercise 7.2: Improving item-based recommendation

Simple item-based recommendation can sometimes be misleading: Many online shops offer accessories for products using item-based recommendations (“other people who bought this lens also bought this filter, this case, this extender, etc”). There is, however, no guarantee that these accessories will match the product, because they could have been bought for another product (another lens in the example). This may lead to unwanted purchases by the customer. Discuss extensions of the recommendation mechanism by which such effects could be reduced, for example by adding external knowledge.

Exercise 7.3: Efficient item-centric recommendations

Design an efficient algorithm for finding the best k items to recommend to a user, given the set of items that he already owns. Hint: precompute, for each item, the similarity to each other item, and store for each item the similar items in descending order of similarity.

Exercise 7.4: Live Evaluation

Design a live user experiment for evaluating the quality of different recommendation algorithms on a shopping site. Assume that the owner of the shopping site wants to identify the best recommendation algorithm for his site, and he is willing to let you modify the site for some time. Discuss alternatives for the design with their pros and cons. Be aware that the main goal of the site is selling items, not evaluating recommender systems, so it is important that users are not overloaded with work. Is it possible to derive the quality of a recommendation algorithm without explicitly asking the user for his opinion?

Exercise 7.5: Anonymity in test collections

For a realistic evaluation of recommender systems, it is necessary to make a large collection of users and their votes for items publicly available, the NetFlix price is a good example for this. These data sets are usually anonymized by replacing user names with random ids. Discuss consequences of this for the evaluation (hint: how can you run a user experiment with this data?). At the same time, it is still possible to break the anonymity of single users by including external knowledge. Discuss possible ways to identify users, especially those with unusual taste (hint: consider information that some users make available on the internet, for example in social networks, about their movie preferences). What could be done in addition to using random user ids to improve anonymity?

Discuss if similar problems can also show up when evaluating approaches for personalized search using query logs, and how they may be overcome (hint: a prominent example of a publicly available query log was made available by AOL some years ago).