Recommender Systems Research at Yahoo! Research Labs

Dennis Decoste*, **David Gleich****, **Tejaswi Kasturi***, **Sathiya Keerthi***, **Omid Madani***, **Seung-Taek Park***, **David M. Pennock***, **Corey Porter***, **Sumit Sanghai****, **Farial Shahnaz****, **Leonid Zhukov***

*Yahoo! Research Labs,
{decosted|tejaswi|keerthis|madani|parkst|pennockd|porterc|zhukovl}@yahoo-inc.com
**Work conducted at Y! RL, dgleich@stanford.edu, sanghai@cs.washington.edu, shahnaz@cs.utk.edu

ABSTRACT

We describe some of the ongoing projects at Yahoo! Research Labs that involve recommender systems. We discuss recommender systems related problems and solutions relevant to Yahoo!'s business.

INTRODUCTION

A number of projects at Yahoo! Research Labs[1] involve collaborative filtering, recommendation, and personalization. A number of business units within Yahoo! either currently use recommendation technology with success, or plan to implement recommendations and personalization in the future. This paper briefly describes some of the relevant projects ongoing at Yahoo! Research Labs.

MAD6

MAD6 (**M**ovies **A**ctors **D**irectors; **6** degrees of separation) is a prototype movie search engine with five major design goals: (1) leveraging relational data implicit in the graph of movies and people (2) extensive metadata indexing, (3) leveraging user ratings and activity logs for personalization and recommendation (4) pseudo natural language query support ("shortcuts on steroids"), and (5) intelligent search ranking based on a combination of popularity, relational inference, and personalization.

We have a working prototype that supports extensive indexing beyond title/name matching (so, for example, queries like "arnold action" and "neo trinity" return meaningful results), intelligent ranking based on popularity and ratings, browsing by movie graph relations, related movie information by overlapping casts and by overlapping user interest, within-genre recommendations and global recommendations. We are currently experimenting with a number of machine learning algorithms to support recommendations, including cold-start recommendations (85% of movies do not have any explicit ratings). We plan to implement other personalizations like allowing a user to search their past activity, to examine plot words, genres, actors, etc. that he/she tends to visit, and to view a "prototypical" movie based on the user's browsing behavior. Finally, we plan to develop a pseudo natural language query interface to MAD6.

[1] `http://research.yahoo.com`

Workshop: Beyond Personalization 2005
IUI'05, January 9, 2005, San Diego, California, USA
http://www.cs.umn.edu/Research/GroupLens/beyond2005

GROUP RECOMMENDATIONS

Group collaborative filtering is the use of individual ratings of various items to form a consensus recommendation for an entire group of people. This can be useful when choosing a form of entertainment for a set of people, such as a movie for a group of friends, or a restaurant for a family. Group recommendations can come through various algorithms, including collapsing a group into a single fake user and using traditional collaborative filtering methods, using a "least misery" method wherein the worst-off person in the group still has acceptable results, or using various voting methods to "elect" an appropriate best choice. We are beginning by tackling the problem of recommending music to a group of friends, building the functionality on top of the open source **Co**llaborative **F**iltering **E**ngine (CoFE) developed at Oregon State University.[2] As in any such work with a subjective output, the difficulty is in finding the best method with which to evaluate the results without requiring a large number of users over a long time period.

ALGORITHMS

Dimensionality reduction

We have employed singular value decomposition (SVD) / latent semantic indexing (LSI) to provide recommended keywords for Overture advertisers to bid on, based on keywords they and others are currently bidding on. We have also employed SVD/LSI for more standard recommendation problems in the movies and music domains.

[2] `http://eecs.oregonstate.edu/iis/CoFE/`

Learning relative/ordinal rankings

Many machine learning methods for recommenders focus on learning numeric ratings. Yet, users are often more comfortable/confident articulating their preferences as relatives (e.g. "like X more than Y") than absolutes (e.g. "like X at level 5 and Y at level 4"). The traditional problem with learning pairwise ranking functions is that this can involve training times that scale quadratically with the number of votes. To overcome this problem, we have been developing fast new methods that, under certain sets of practical situations (such as linear models) scale only log-linearly.

Content vs. collaborative filtering

In many instances we have extensive metadata about items, and sometimes demographic data about users. We are exploring a variety of machine learning algorithms to mimic collaborative similarities using content data, which can then be applied to sparse regions of the data.

Active learning and the cold-start problem

One key problem is what to do with new users (often a large fraction of users) and new items. We are exploring active learning techniques to determine which questions to ask users. For example, in the movie domain, we can ask for ratings on movies, actors, directors, genres, awards, etc. We can ask for numeric ratings or comparisons. We can ask more generic personality questions ("do you cry at movies?", "which among this set of abstract images[3] appeals to you?"). Choosing the right question involves determining informativeness, ability to answer, and willingness to take the time to answer (user burden). Determining the best questions is ultimately an exploration/exploitation tradeoff, since we won't know a question's informativeness until we receive sufficient answers on which to base inferences.

Music recommendation via audio similarity

With the recent explosion of availability of MP3 and other music data sources, significant recent research activity has focussed on new methods for summarizing music audio data and defining suitable similarity measures to supplement traditional metadata and ratings data. Often the somewhat hard-to-predefine audio nature of a song (i.e. combination of beat style, existence of unique guitar riffs, overall audio "feel", etc.) has more impact on whether someone likes the song than on what genre or artist is associated with it. If asked, most everyone seems to think that their own musical tastes are "eclectic"—hard to characterize as some simple cluster in artist/genre space and not necessarily particularly similar any other user. Exploiting such audio content presents huge new machine learning challenges, due to the significantly increased raw dimensionality (i.e. millions of bits of raw audio data per song) of the content data, and determining similarity metrics that correspond as closely as possible to human perception.

[3] `http://www.cs.ucr.edu/~chua/`

INTERFACES

We have built a prototype "World of Music" searchable map, which is a low-dimensional projection of music artists proximity for visual display of artist-artist similarity. We have also built a Java applet that implements a simple spring-force-based layout for exploring the space of recommendations and related items in the movie and music domains. We plan to use MAD6 as a platform for testing various search, browsing, personalization, and recommendation interfaces in the movie domain. Some of these demos may become available in the future via the lab website, `http://research.yahoo.com`.

EVALUATION METRICS

We plan to explore the utility of several offline and online metrics, with the goal of determining which offline metrics best predict important online metrics. As a company with a large user base, it is possible to try beta algorithms on small percentages of traffic and still obtain meaningful statistics. Moreover, as a largely advertiser-funded media company, Yahoo! can mainly focus on satisfying users to encourage retention, without need to consider inventory for example.

YAHOO! BUSINESS APPLICATIONS

A number of Yahoo! properties and business units use recommendation technology, or are planning or considering using recommendation technology, including Launch Music on Yahoo!, MusicMatch, Yahoo! Movies, Yahoo! TV, Yahoo! Personals, Yahoo! Local, Yahoo! Autos, Yahoo! Search, targeted banner advertising, Overture sponsored search advertising, and Overture contextual advertising. A company-wide effort to offer packaged recommendation technology to any interested Yahoo! property is underway.

ACADEMIC COLLABORATION

We try to maintain close ties to the academic community by staying current on the latest research, publishing our own research results, attending and sponsoring relevant workshops and conferences, hiring graduate student interns, hosting faculty sabbaticals, and hosting *spot workshops* on site. For example, in August 2004, we hosted a spot workshop on recommender systems featuring both external academic speakers and internal Yahoo! speakers from various business units.[4]

We have had success using and building on Oregon State's CoFE software. We have been able to share data on a case by case basis with academic collaborators, and would like to expand our data sharing activities.

ACKNOWLEDGMENTS

We thank Todd Beaupre, Donna Boyer, Nilesh Gohel, and Steve Souders.

[4] `http://research.yahoo.com/~pennockd/spot/rs/`