

CSCW 2000 Report

Recommender Systems

By Andy Gorman

Flavors of Recommender Systems

I attended the tutorial, “Recommender Systems: Collaborating in Commerce and Communities” at this year’s CSCW conference. Recommender systems attempt to alleviate the problem of information overload. There are too many books, journal articles, news articles, etc. for us to sift through. Recommender systems address this problem by recommending things that are “worthwhile”. There are many approaches used in supplying recommendations. Some simply supply the opinions of supposed experts, such as professional movie critics. One of the problems is that what one critic considers worthwhile another might consider as a waste of time. Some people have learned that a certain critic represent their personal tastes better than others, but many still settle for the “one size fits all” solution.

Collaborative Filtering

Many recommender systems uses some form of social or collaborative filtering. An example of this is MovieLens (<http://movielens.umn.edu>). MovieLens uses the reviews of its users to make recommendations. By comparing one user’s taste in movies to other users in the system, MovieLens tries to make recommendations. It accomplishes this by finding a user’s “neighbors”. These are people that have rated movies most similarly to the user in question. It then recommends movies that the user has not yet rated but that were rated highly by their neighbors (Herlocker, Konstan, Borchers, & Riedl, 1999). This type of approach does something that a professional movie critic can not do: it accounts for the diverse and subjective tastes of the individual community members.

Explicit vs. Implicit Recommendations: Reducing User Effort

In the previous case, a user’s preferences are explicitly expressed (i.e., a user explicitly rates movies that he or she has seen). This places quite a burden on users A user must rate at least 5 movies before MovieLens can attempt a prediction. Other systems, such as PHOAKS (People Helping One Another Know Stuff, <http://phoaks.com>), use web-mining techniques to extract more implicit recommendations. PHOAKS scans Usenet news groups for messages that contain references to other web sites. Using semantic analysis techniques, it determines whether the reference is an endorsement or not. It then uses the frequency of endorsements for a web site to rank (and therefore recommend) web sites. The popular search engine, Google.com uses a similar approach It considers one web site to be more authoritative than another site if a greater number of other web sites link to it. The benefit of automatically extracting

implicit recommendations is that there is no effort required of the user. It should be noted, however, that while all these systems (MovieLens, PHOAKS, and Google) use a form of social filter, only MovieLens takes the personal tastes of its users into account.

Past Usage as a Way of Predicting Future Needs

Some systems analyze users' behavior or usage patterns and then make predictions based on those patterns. For example, Amazon.com keeps track of what each user has purchased, viewed, and rated. Based on this information it makes recommendations on what a user might want to consider next.

Hybrid Systems

As this field becomes more and more mature, developers are beginning to create systems that combine various approaches in a way that produces better results than any single approach.

CLever Recommender Systems

I suppose that it's true that when you have a hammer, everything looks like a nail (i.e., now that recommender systems are fresh in my mind, I tend to see it as the solution to many of our problems). Below I describe two possible applications of recommender systems in the CLever project.

Expertise Recommender (ER)

David McDonald and Mark Ackerman from UC Irvine discussed their experience with ER (McDonald & Ackerman, 2000), a system that helps people find experts in a certain problem domain. A system like this can be useful for teachers and parents alike. Several special ed. teachers have noted that sharing expertise has been a problem. Consider a teacher in a rural community that has one student with special needs. There are no local resources for this teacher so how does he or she learn how to help their student? I think that a system like ER might help elevate this problem.

As I noted above, I think that such a system can be useful for both teachers and parents. Whether these two populations can be integrated into one support system or not is a question that needs to be discussed. Irrespective of the answer, I think that "Experiences Journals" or "Life Stories" can help in matching up people. I briefly discuss this in another document titled "CATP 2000 Report" (see <http://www.cs.colorado.edu/~l3d/clever/pdf/ag-200011-catp-report.pdf>), where I describe the Parent to Parent of Colorado organization. It should be noted that there seems to be a link between what Parent to Parent and ER tries to accomplish.

One of the interesting facets of ER is that it recognizes the need to account for the social idiosyncrasies of a group and its members. Just because someone is an expert in a certain area doesn't mean that they are always willing or able to share their expertise. Realizing that an individual's status fluctuates over time is an

important component to expertise recommendations. While ER was designed to support a software development group with well-defined work practices, the paper they presented at the conference describes a more general-purpose architecture for building such systems. I think this will be a useful paper to study in more detail.

Assistive Technology Recommendations

After seeing the MovieLens project, I began to wonder whether assistive technology could be recommended in a way similar to movies. I immediately began to see problems. The nature of the objects being evaluated is very different. A movie is released, and from that point on it is a static entity—it does not change. Assistive technology, on the other hand, becomes outdated and replaced by newer technology over time. A challenge in an assistive technology recommendation system will be to properly account for the temporal dimension. There needs to be a way of linking assistive technology products across time. For example, product *x* worked well in the '80s, product *y* is the most similar and current product on the market. We also need to look at whether we want to attempt to recommend specific products or, more generally, a family of products (e.g., I recommend a portable communication devices that use an iconic language system vs. I recommend a Dynavox model 2300).

A second difference between assistive technology and movies is in the usage patterns. A typical user of movie recommendations has seen dozens of movies. In fact, MovieLens relies on this usage pattern: the more movies a user rates, the better the systems predictions become. Someone interested in an assistive technology recommendation may very well be a first-time buyer. I think that “life stories” can be useful in overcoming this problem. Instead of trying to find similarities in ratings, as is done in MovieLens, we may be able to look for similarities in life stories (i.e., “My situations is similar to theirs, what have they used and liked?”). A system that uses case comparison can also be supplemented with a categorization schema or taxonomy developed by health sciences. In other words, the bottom-up case-based approach can be guided by a top-down classification approach.

Integration and reuse of information

In this document I mentioned only two possible uses of recommender technology. For simplicity's sake I described them somewhat as isolated systems. I do not believe, however, that this is necessarily the best approach. For example, the output of a system that links parents together can also be used as input to a system that recommends assistive technology. For example, if I am looking for technology for my child, a first step might be to find other parents with children in similar circumstances. Then I can look at what technology has worked for them and base my decision on that information. On the other hand, I might just want to find parents with circumstances similar to mine so that I can establish a relationship with them. Therefore, it is my opinion that there should be separate systems: one for linking parents and one for recommending technology. In this case the former can guide the later.

Similarly, a disability taxonomy and experience journal can be used as input for both systems. It makes sense that as our project evolves, we should look for these opportunities to share data and create synergies.

References

Herlocker, J. L., Konstan, J. A., Borchers, A., & Riedl, J. (1999). *An algorithmic framework for performing collaborative filtering*. Paper presented at the 22nd annual international ACM SIGIR conference on Research and development in information retrieval, Berkeley, CA USA.

McDonald, D. W., & Ackerman, M. S. (2000). *Expertise Recommendation: A Flexible Recommendation System and Architecture*. Paper presented at the Proceedings of the ACM 2000 conference on Computer supported cooperative work, Philadelphia, PA, USA.