

Putting Recommender Systems to Work for Organizations

Natalie S. Glance

Xerox Research Centre Europe
6 chemin de Maupertuis
38240 Meylan, France
glance@xrce.xerox.com

Introduction

What recommender systems have in common is an emphasis on leveraging social processes for the purpose of improving information access. Typically, most of the current breed of recommender systems are Internet services with a two-fold purpose: providing tailored recommendations and building communities.

The issue we focus on here is how to make recommender systems work *in* organizations and *for* organizations. Moving from the Internet to Intranets requires shifting the primary focus from sharing recommendations to sharing knowledge and from community-building to community support. Moving recommender systems from the Internet onto Intranets also means turning "leisure-ware" into groupware, creating both new challenges and new opportunities.

Research issues

When we think about putting recommender systems to work for organizations, we need to re-think a number of assumptions underlying Internet-based recommender systems: for example, the potential size of the user base (huge and still rising exponentially) and the *raison d'être* (recommendations oriented towards leisure and personal interests). In contrast, in work organizations, the potential user base will be relatively fixed and of more modest size. Also, the main impetus, or at least, the major impetus for usage will be sharing information pertinent to workplace competencies, interests, and goals.

On the Internet, a useful service can hope for a user base of thousands, if not millions, of users. Recommender systems that depend on statistical algorithms thrive on extensive usage. In contrast, recommender systems in primarily closed organizations of more modest size must implement new ways to make recommendations from a smaller user base. Many Internet-based systems using collaborative filtering techniques have already discovered the benefit of filtering first by content and then by taste. This two-tiered approach is likely to be even more important in work environments where differences in preferences cannot be accounted for simply by taking into

account the opinions of yet more people rating yet more items.

However, filtering first by content and then by taste means calculating correlations based on far fewer rating pairs, a problem which is likely to be far worse in an organizational setting where interest domains can be populated by mere handfuls of people. These problems highlight the importance of going beyond automated collaborative filtering for making recommendation predictions. To make good predictions, organization-based recommender systems will have to take into account usage data across many systems, potentially including search engines, document management systems, possibly e-mail, and will have to finesse methods for combining the many kinds of evaluations into one form.

Recommender systems as groupware instead of leisure-ware also suffer much more from the critical mass problem. Within a work organization, a recommender system can be a valuable tool only if most (as opposed to many) people are using it. Like the telephone or the fax or e-mail, this kind of technology will be used at work only if most others use it as well. This is true of Internet-based recommender systems, but to a lesser extent, as the Internet can make up in numbers what it lacks in density.

Thus, it is vital in a work setting to investigate the incentive issues that arise in achieving acceptance and usage of the technology. These incentives include: ease-of-use, integration with users' software environment, and perhaps most importantly, immediate and sustained perceived benefit. This last issue is particularly difficult: how to overcome the well-known cold-start problem of recommender systems to provide useful and high-quality recommendations immediately?

Perhaps most interestingly, we can expect the dynamics of group participation and usage to be qualitatively different in a work situation where people know each other by name and reputation. How can work-based recommender systems leverage the notions of trust, reputation and reciprocity to best serve working communities and help them serve themselves? Here, we can expect to learn both from other peer-based models for collaborative work and evaluation and from market mechanisms for privatizing public goods (evaluations being a public good). However, the true test will be in the usage:

in this sense, work-centric recommender systems are just like their Internet cousins; they serve as platforms for novel social experiments and institutions.

Finally, while Internet-based recommender systems focus on *creating* communities by bringing people together, Intranet-based recommender systems should focus on *supporting* communities that already exist. The power of recommender systems to help people find each other is multiplied many times over in an organization: locating experts; reducing re-work by bringing together people working in the same area in different geographic locations; and identifying competencies, both established ones and emerging ones.

Current research activities

At XRCE Grenoble, we are pursuing several concurrent threads of activity regarding organization-based recommender systems:

- We are developing a recommender system component as part of a larger knowledge management platform for supporting knowledge sharing.
- We are investigating new prediction algorithms

enhancing content-based collaborative filtering that (1) are able to immediately make reasonably good recommendations for a new user; and (2) take into account usage data across different software tools.

- We are studying how system feedback, visualization tools and market mechanisms can be used to influence the dynamics of group participation.

Currently, we have a working research prototype of our knowledge sharing system that is being used by a group of 30 people internally (see Figure 1). In the short term, we will be enlarging the user base to include the Grenoble lab and affiliates within Xerox.

Further reading

Glance, N. S., Arregui, D., and Dardenne, M. 1998. Knowledge Pump: Supporting the Flow and Use of Knowledge. In *Information Technology for Knowledge Management*, Ed. U. Borghoff and R. Pareschi, NY: Springer-Verlag.

Glance, N. S., Arregui, D., and Dardenne, M. 1998. Making Recommender Systems Work for Organizations. Submitted to CSCW'98.

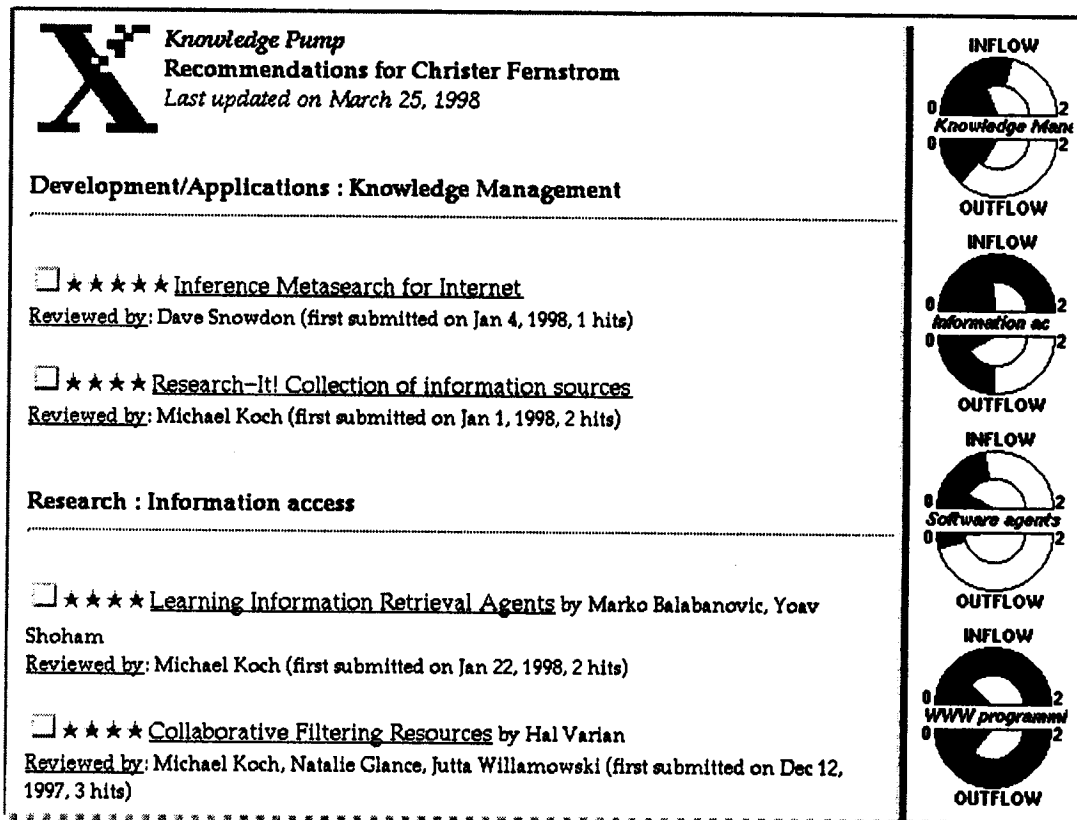


Figure 1. "What's Recommended?" by the Knowledge Pump today.