

# *Who's Who in Recommender Systems*

Non-attendees

## **Amazon.com**

"Our goal is to provide our 20 million users with 20 million different stores," Jeff Bezos explained at PC Expo 2000. "We're the most customer-centric company on earth."

Amazon.com personalization technology records what a shopper purchases, rates, or views on the website and then matches that data to the aggregated purchases of other shoppers who have bought similar products. By comparing what an individual shopper has purchased to aggregates of what shoppers with similar tastes have purchased, Amazon.com can provide highly targeted, relevant, and interesting recommendations to our customers. Features that use this technology on Amazon.com's website include New for You, Instant Recommendations, Gift Wizard, Recommendations Wizard, Similarity Explorer, Page You Made, "Customers who bought...", and personalized new releases.

"I want to transport online bookselling back to the days of the small bookseller, who got to know you very well and would say things like, 'I know you like John Irving, and guess what, here's this new author, I think he's a lot like John Irving.'", said Jeff Bezos. Of course, in Amazon.com's case, it's the technology, and not the antiquarian in the musty store, that remembers your likes and dislikes.

Amazon.com ([www.amazon.com](http://www.amazon.com)) opened its virtual doors on the World Wide Web in July 1995 and today offers Earth's Biggest Selection, along with online auctions and free electronic greeting cards. Amazon.com seeks to be the world's most customer-centric company, where customers can find and discover anything they might want to buy online. Amazon.com lists millions of unique items in categories such as electronics, kitchen and housewares, books, music, DVDs, videos, camera and photo items, toys, software, computer and video games, tools and hardware, outdoor living and wireless products. Amazon.com operates four international Web sites: [www.amazon.fr](http://www.amazon.fr), [www.amazon.co.uk](http://www.amazon.co.uk), [www.amazon.de](http://www.amazon.de) and [www.amazon.co.jp](http://www.amazon.co.jp).

# Recommender systems: Interfaces and Architectures

Marko Balabanović

flutter.com

175-179 St John St, London EC1V 4LW, UK

marko@bigfoot.com

## INTRODUCTION

This paper provides a brief summary of research into interfaces and architectures for recommender systems, undertaken during 1994-98. A complete description can be found in [3].

We consider specifically the text recommendation task: helping the user follow threads of interest by regularly recommending small sets of 6-10 documents. Systems select a set of documents to present to the user, based on an induced user model; elicit feedback from the user (either explicit or implicit); and finally improve the user model given the feedback on the presented documents.

## RESEARCH ISSUES

Two specific research issues are addressed by the research described here.

First, how to exploit the overlap between users' interests. In a system with multiple users, each with multiple topics of interest, resources can be allocated to serve areas where different users' interests overlap. Users' feedback can also be pooled—a collaborative strategy that can improve recommendations as well as reducing duplicated requests for user feedback.

Second, how to decide upon the composition of recommendation sets. Techniques for choosing individual articles for recommendation are well known, but another important consideration is how to compose a set of such recommendations (e.g., the selection of articles for an edition of a personalized newspaper).

## COMBINING CONTENT-BASED AND COLLABORATIVE RECOMMENDATION

Content-based recommendation requires analysis of the content of items. This poses several problems: only certain domains can be readily analyzed (text is much easier than music or restaurants); systems will often over-specialize, only recommending items similar to those already rated highly. Collaborative recommendation instead computes similarity between users. As pointed out in [1], it avoids the problems of a pure content-based system. However it has its own shortcomings: new items cannot be recommended until some user has rated them, making very dynamic domains difficult; users with unusual tastes will be served poorly; negative feedback about items will not guarantee similar items won't be recommended again.

The *Fab* system combined content-based and collaborative recommendation to eliminate disadvantages of using either approach alone. A pool of *collection agents* adaptively gathered Web pages matching topic

profiles. Each user's *selection agent* learned a personal user model and composed recommendation sets from the discoveries of the various collection agents and from other selection agents with similar models. The collection agents thus converged on areas of overlap between the users' interests, and the selection agents served both content-based and collaborative recommendations.

## EXPLORATION vs. EXPLOITATION

A large-scale simulation was constructed in order to study the effects of varying the "breadth" of a set of recommendations [2]: whether to recommend items similar to those that have already received positive feedback, or new items where the user's reaction cannot be predicted. Experiments showed that an exploratory strategy increased the rate of learning user profiles, and the rate of adaptation to changing user interests, at the expense of showing users documents that are less often to do with their preferred topics. The result of this research was a new affordance for the *Fab* system allowing users to navigate among "broad" and "narrow" recommendation sets.

## INTERFACES THAT CAPTURE IMPLICIT FEEDBACK

The *Slider* system [4] took these ideas further. Overlapped, sliding panels allowed users to specify multiple topics of interest and interactively alter the proportions between them. In addition, the interface did not require explicit feedback from a user, but successfully learned user profiles by monitoring drag-and-drop, selection and deletion actions. An "other news" topic was always available, showing items with the least overlap to those previously recommended.

## REFERENCES

1. Balabanović, M., and Shoham, Y. *Fab: Content-based, collaborative recommendation*. *Commun. ACM* 40, 3 (1997).
2. Balabanović, M. Exploring vs. exploiting when learning user models for text recommendation. *User Modeling and User-Adapted Interaction* 8, 1 (1998).
3. Balabanović, M. Learning to Surf: Multiagent Systems for Adaptive Web Page Recommendation. Stanford University Department of Computer Science, PhD Thesis (1998).
4. Balabanović, M. An Interface for Learning Multi-topic User Profiles from Implicit Feedback. *Proc. of AAAI Workshop on Recommender Systems* (Madison WI, 1998), AAAI Press.

# Research in Recommender Systems at Worcester Polytechnic Institute

Lee Becker, David Brown, Mark Claypool and Carolina Ruiz

<http://www.cs.wpi.edu/Research/refer/>

REsearch in FiltEring and Recommender systems (REFER) is a research group in the Computer Science department at Worcester Polytechnic Institute, in Worcester, Massachusetts, USA. The members of REFER share interests in issues relevant to the design and analysis of personalized information systems that recommend items of potential interest to their users on the basis of descriptions of these items (“content”) as well as social (“collaborative”) information about the relations between different users’ tastes.

To provide personalized recommendations a recommender system must have ratings on each item from each user. The common solution, *explicit ratings*, where users tell the system what they think about a piece of information, is well-understood. However, stopping to enter explicit ratings can alter normal patterns of browsing and reading. Using *implicit ratings*, obtained by a method other than obtaining it directly from the user, has obvious advantages, including removing the burden on the user, and that every user-system interaction can contribute to an implicit rating. A main objective of our research is to collect, measure and evaluate the predictive power of implicit interest indicators, as begun in [CLWB01]. We are in the process of evaluating a large number of indicators, such as visiting a page, bookmarking, mouse movement, printing and time spent scrolling, and combinations of these indicators in order to determine which can be used to reliably predict interest in a Web page, and which might be combined. Through experiments, user studies and data analysis, we intend to show how implicit interest indicators can predict explicit interest, for all users, for groups of users, for an individual user, by using both individual and grouped implicit indicators.

We are also investigating the use of association rule mining as an underlying technology for collaborative recommender systems, as in [LAR02]. Association rules have been used with success in other domains. However, most currently existing association rule mining algorithms were designed with market basket analysis in mind. Such algorithms are inefficient for collaborative recommendation because they mine many rules that are not relevant to a given user. Also, it is necessary to specify the minimum support of the mined rules in advance, often leading to either too many or too few rules; this negatively impacts the performance of the overall system. We have developed a collaborative recommendation technique based on a new algorithm specifically designed to mine association rules for this purpose. This algorithm does not require the minimum support to be specified in advance. Rather, a target range is given for the number of rules, and the algorithm adjusts the minimum support for each user in order to obtain a ruleset whose size is in the desired range. Rules are mined for a specific target user, reducing the time required for the mining process. We employ associations between users as well as associations between items in making recommendations. Experimental evaluation of a system based on this algorithm reveals performance that is significantly better than that of traditional correlation-based approaches.

While collaborative recommendation combines the informed opinions of humans to make personalized, accurate predictions, content-based recommendation uses the speed of computers to make complete, fast predictions. Our research also explores new filtering approaches that combine the coverage and speed of content-based recommendation with the depth of collaborative recommendation, as in [CGM+99]. Our research also extends our earlier association rule approach [LAR02] to handle content-based recommendation, and a combination of content-based and collaborative recommendation.

[CGM+99] M. Claypool, A. Gokhale, T. Miranda, P. Murnikov, D. Netes and M. Sartin, “Combining Content-Based and Collaborative Filters in an Online Newspaper”, In *Proceedings of ACM SIGIR Workshop on Recommender Systems*, Berkeley, California, USA, August 1999

[CLWB01] M. Claypool, P. Le, M. Waseda and D. Brown. “Implicit Interest Indicators”, In *Proceedings of ACM Intelligent User Interfaces Conference (IUI)*, Santa Fe, New Mexico, USA January 2001

[LAR02] W. Lin, S.A. Alvarez and C. Ruiz, “Efficient Adaptive-Support Association Rule Mining for Recommender Systems”, *Data Mining and Knowledge Discovery Journal*, January 2002 (To Appear)

For the past year Britannica.com, Inc. has been conducting extensive analysis of the web usage logs from the Britannica.com website to improve various characteristics of the site. At this time, researchers at Britannica are hoping to take this analysis further to begin recommending encyclopedia articles, news stories, products, etc. to the users of the site based on past and current user behaviors.

Since this work is in the infantile stage, many aspects have not yet been defined. It is expected, though, that Britannica researchers will take into account user clicks and queries in an attempt to identify related concepts. These related concepts will be used to recommend web pages and products that may be of interest to a current user.

Researchers involved with this work:

Laura Mather

[lmather@us.britannica.com](mailto:lmather@us.britannica.com)

Kunal Sen

[ksen@us.britannica.com](mailto:ksen@us.britannica.com)

Britannica's website:

[www.britannica.com](http://www.britannica.com)

## Recommender Systems at the University of Glasgow

Matthew Chalmers

[www.dcs.gla.ac.uk/~matthew](http://www.dcs.gla.ac.uk/~matthew), [matthew@dcs.gla.ac.uk](mailto:matthew@dcs.gla.ac.uk)

Our work spans theory and system development, and explores the use of a recommender system within multiple media and tools. It builds on earlier experimentation with a contextually-specific recommender for the Web and local file systems, *Recer* (Chalmers et al. 1998). *Recer* makes recommendations by comparing each person's ongoing context with his or her past activity and/or the past activities of others. Past activity here is represented by a time-stamped log of locations, documents, artefacts and people that each person has interacted with in all of the media we can: that person's *path*. Unlike most recommender systems, recommendations are specific to the ongoing context of the user, draw upon only the most relevant periods of past activity, and combine physical locations, hypermedia 'locations' and references, and locations in virtual environments. As discussed in (Chalmers 1999) and (Chalmers 2001), this approach is based on a combination of aspects of urban design theory, the trails of Bush's *As We May Think*, linguistics, semiology and philosophy of language.

Most of our work is centred in the *City* project. While we will gradually extend the range of places and topics it handles, the initial focal point is Charles Rennie Mackintosh, the architect, designer and artist. A rich body of 'people, places and things' related to Mackintosh exists here in Glasgow, including his reconstructed house within the Hunterian Gallery and Museum, the exhibition room devoted to his life and work within the Lighthouse Centre, and buildings he designed, e.g. the Art School and the Lighthouse itself. The Hunterian and the Lighthouse have agreed to be partners and test sites in this project, and in other related projects within the six year, eight site and £11,000,000 consortium called 'Equator'. Equator ([www.equator.ac.uk](http://www.equator.ac.uk)) is an 'interdisciplinary research collaboration' funded by the UK's Engineering & Physical Sciences Research Council and involving computer scientists, sociologists, psychologists, artists and designers. University College London, the Royal College of Art and the universities of Bristol, Glasgow, Nottingham and Southampton are taking part in the *City* project, with Glasgow as lead.

We are combining mobile computers, hypermedia and virtual environments in one system, and tracking each person's activity in them all as a path. We allow each person to interact with others even if they are using quite different media or combinations of media. We have found it useful to consider the many media, technologies and spaces as one design medium, because each person's experience depends on them all. People's activity continually combines and cuts across different media, interweaving those media and building up the patterns of association and use that make meaning. How people act and work is determined by the full combination of media that they can use and have used, and hence a narrow focus on only digital media as the paramount determinant of activity underrates the influence of other media. Irrespective of the media involved in logging and display, our algorithms for searching, matching and recommending do not rely on distinctions between these original media. They rely on their patterns of co-occurrence in human activity and hence in common semiological use. It is this that lets recommendations bridge across media usually held as separate in information systems, and lets us work with this everyday mix of types rather than against it. For example, a person who has been looking at the original chairs and tables made by Mackintosh may get recommendations of other physical artefacts to look at, but also of digital documents and virtual locations that might be of interest. These recommendations may come from people who have never seen the physical furniture, but have explored related digital information. One is offered information based on one's current location and the route one has recently taken, the information read and written, the artefacts one has shown an interest in—and how this activity relates to the activity of earlier readers, authors and visitors. The past routes and paths of curators, designers, authors and visitors are combined with current context to suggest recommendations for the immediate future.

We are creating a growing and evolving body of individuals' paths or narratives through the people, places and artefacts associated with the city of Glasgow. An initially static collection includes images, textual descriptions, and references to locations in the city, to artefacts in museums and exhibitions, and to electronic resources such as web pages and virtual 3D environments. Later, we will allow this body of information to grow as people use it, with their paths making new associations between symbols and also adding in new symbols. *City* includes information access based on objective bases, such as maps and search engines, and subjective bases, such as recommenders and synchronous audio links. We have now integrated VR and PDA access to the Mackintosh Room, will integrate a hypermedia interface within the next month, and the recommender within two months. We plan to demonstrate this prototype by October 2001, with synchronous and asynchronous awareness across all three media supported, paths interweaving symbols from all media and used as a resource for recommendations, and first experiences of non-Equator people to report and discuss. We have recently begun a series of semi-structured interviews with curators and exhibition designers in a number of UK museums, and a sociological study of related activity by city visitors and residents. These studies will feed reflection on our first full prototype and the design of future systems.

[Chalmers 1998] Chalmers, M., K. Rodden & D. Brodbeck: The Order of Things: Activity-Centred Information Access, In *Proceedings of the 7<sup>th</sup> World Wide Web Conference*, Brisbane.

[Chalmers 1999] Chalmers, M.: Informatics, Architecture and Language. In A. Munro, K. Höök & D. Benyon (eds.): *Social Navigation in Information Space*, Springer, 1999.

[Chalmers 2001] Chalmers, M.: Awareness, Representation and Interpretation, to appear in *J. Computer Supported Cooperative Work*.



# Web Link and Traffic Analysis for Recommendation

Brian D. Davison

Department of Computer Science and Engineering

Lehigh University

davison@acm.org

The Web has two pervasive sources of activity that can be useful for recommendation systems. One is the traffic on the Web, reflected in individual user histories as well as proxy and origin server usage logs. The second is in hypertext link authoring. Both provide information that can often be treated as a value judgment. A user who clicks on a link, in effect, believes that there is some content of value at the destination. The significance of link creation is thus even stronger, as it reflects the judgment of the author (presumably) after viewing the target content and assessing it to be of sufficient value to spend the time to make the connection available to readers.

We have examined the relationship between the description of a page and its contents [4], and explored the value of link analysis for search engine result ranking and community discovery [8], as well as tested for nepotistic links that may influence such analysis [3].

Perhaps unlike most researchers in recommender systems, our interest is not only recommendation for human usage, but also recommendation for systems. Adaptive systems can then use those recommendations (from whatever source) to improve performance invisibly (unlike, for example Letizia [10] or WebWatcher [9] which alter the Web interface that is presented to the user). We are also concerned with appropriate evaluation of such systems.

While in the past we have examined techniques to build user models for user interfaces (such as the UNIX shell [7]), in recent years we have concentrated on issues related to modeling Web users for pre-loading of Web content to improve user-perceived performance [11]. We have proposed adaptive Web prefetching [1] which incorporates prediction from multiple sources, including:

- individual user histories,
- workgroup histories from proxy cache hints (a form of collaborative filtering),
- global user histories from origin server hints (again collaborative filtering),
- predictions from analysis of recently viewed content [6], and
- predictions from other available context, such as bookmarks, e-mail and usenet news.

Evaluation of Web pre-loading systems can be difficult. To that end, we have proposed [2] and implemented one

approach to evaluate prefetching proxy caches simultaneously, and built and validated [5] a detailed network and caching simulator to estimate user latencies for testing and comparing alternate history-based prefetching approaches.

## References

- [1] B. D. Davison. Adaptive Web Prefetching. In *Proc. of the 2nd Workshop on Adaptive Systems and User Modeling on the WWW*, pp. 105–106, Toronto, May 1999.
- [2] B. D. Davison. Simultaneous Proxy Evaluation. In *Proc. of the 4th Int'l Web Caching Workshop (WCW99)*, pp. 170–178, San Diego, Mar. 1999.
- [3] B. D. Davison. Recognizing Nepotistic Links on the Web. In *Proc. of Workshop on Artificial Intelligence for Web Search*, pp. 23–28. AAAI Tech. Rep. WS-00-01, July 2000.
- [4] B. D. Davison. Topical Locality in the Web. In *Proc. of the 23rd Annual Int'l Conf. on Research and Development in Information Retrieval (SIGIR 2000)*, Athens, July 2000.
- [5] B. D. Davison. HTTP simulator validation using real measurements: A case study. In *Proc. of MASCOTS*, Cincinnati, OH, Aug. 2001.
- [6] B. D. Davison. Predicting Web actions from HTML content. Under submission, 2001.
- [7] B. D. Davison and H. Hirsh. Predicting sequences of user actions. In *Proc. of Workshop Predicting the Future: AI Approaches to Time-Series Problems*, pp. 5–12, Madison, WI, AAAI Tech. Report WS-98-07, July 1998.
- [8] B. D. Davison, A. Gerasoulis, K. Kleisouris, Y. Lu, H. Seo, W. Wang, and B. Wu. DiscoWeb: Applying Link Analysis to Web Search. In Poster proc. of the 8th Int'l World Wide Web Conf., pp. 148–149, Toronto, May 1999.
- [9] T. Joachims, D. Freitag, and T. Mitchell. WebWatcher: A Tour Guide for the World Wide Web. In *Proc. of the 15th Int'l Joint Conf. on Artificial Intelligence*, pp. 770–775. Morgan Kaufmann, Aug. 1997.
- [10] H. Lieberman. Letizia: An Agent That Assists Web Browsing. In *Proc. of the Int'l Joint Conf. on Artificial Intelligence (IJCAI'95)*, pp. 924–929, Montreal, Aug. 1995.
- [11] V. N. Padmanabhan and J. C. Mogul. Using predictive prefetching to improve World Wide Web latency. *Computer Communication Review*, 26(3):22–36, July 1996.

## Eigentaste: A Constant Time Collaborative Filtering Algorithm

Ken Goldberg, Theresa Roeder, Dhruv Gupta, and Chris Perkins  
IEOR and EECS Departments  
University of California, Berkeley  
goldberg@ieor.berkeley.edu

Eigentaste is a collaborative filtering algorithm that uses *universal queries* to elicit real-valued user ratings on a common set of items and applies principal component analysis (PCA) to the resulting dense subset of the ratings matrix. PCA facilitates dimensionality reduction for offline clustering of users and rapid computation of recommendations. For a database of  $n$  users, standard nearest-neighbor techniques require  $O(n)$  processing time to compute recommendations, whereas Eigentaste requires  $O(1)$  (constant) time. We compare Eigentaste to alternative algorithms using data from *Jester*, an online joke recommending system.

Jester has collected approximately 2,500,000 ratings from 57,000 users. We use the Normalized Mean Absolute Error (NMAE) measure to compare performance of different algorithms. In the Appendix we use Uniform and Normal distribution models to derive analytic estimates of NMAE when predictions are random. On the Jester dataset, Eigentaste computes recommendations two orders of magnitude faster with no loss of accuracy.

(The Jester dataset including ratings from approximately 18,000 anonymous users is available by request: contact goldberg@ieor.berkeley.edu with contact information and a description of intended research.)

The joke recommending system *Jester* is available online at:  
<http://eigentaste.berkeley.edu/>

The book recommending system *Sleeper* is available online at:  
[www.pmetrics.com](http://www.pmetrics.com)

A .pdf version of our *Information Retrieval Journal* paper (July 2001) is available at:  
[www.ieor.berkeley.edu/~goldberg/pubs/](http://www.ieor.berkeley.edu/~goldberg/pubs/)

**Sukumal Imudom**

[imudom@usc.edu](mailto:imudom@usc.edu)

Computer Science Department  
University of Southern California  
University Park Campus  
Los Angeles, CA 90089

**Summary**

I have been working on distributed annotation framework supporting collaborative filtering of any type of information. Currently, there exists the problem of validity and trust in recommendations placed on such systems. The framework provides a mechanism to build trust and reputation of the evaluators, as well as a reward model as an incentive for promoting evaluations.

The framework consists of three important building blocks aiming to overcome existing recommender system limitations: annotation semantic categories, expertise rating of the annotators, and a credit transfer mechanism. A generic set of annotation semantic categories is defined in order to help classify the free text annotations and provide annotation structure. The notion of "expertise" is introduced in the framework to help identify experts in our system. The expertise level of evaluators and a mechanism to identify the usefulness of evaluations can be used to help determine the credibility of the information. Additionally, we propose a credit transfer model to overcome the cold-start problem which is the lack of initial evaluations by providing incentives to promote user participation. This model, a hybrid of the subscription and transaction models, uses electronic credits as rewards for evaluators who make early and useful annotations.

The framework acts as a middleware allowing application developers to build various applications benefiting from the filtering process on top of it. The future use of these applications in the integration of recommender systems and Electronic Commerce is envisaged. For example, evaluators with credibility can earn community rewards. Such rewards can be mapped to money while giving their evaluations to customers. Not only would the future recommender systems help users more efficiently filter the quality of information, it also would provide a new way for consultants or analysts to do their jobs.

**References**

Sukumal Imudom and B. Clifford Neuman, "**A Framework Supporting Collaborative Filtering for Internet Information**", AAAI-98 Workshop on Recommender Systems, Madison, Wisconsin, July 1998.

## The Active WebMuseum

Arnd Kohrs, Prof. Bernard Merialdo  
Multimedia Communications Department  
Institut EURECOM  
BP 193, 06904 Sophia-Antipolis, FRANCE  
{kohrs,merialdo}@eurecom.fr  
<http://bmgroup.eurecom.fr/>


### Project Summary

Our global interest lies in Information Filtering. In the Active Web Museum project, we are studying several issues related to the usage of collaborative filtering techniques :

- How to cope with the problem of insufficient data, in order to get reasonable performance even when little data is available,
- How to mix collaborative and content-based prediction,
- How to use prediction to personalize a web site.

The Active Web Museum is a web site which uses both collaborative and content-based techniques to predict paintings of interest for the user. Paintings are organized along virtual corridors, and the prediction value is used to order paintings inside the corridor. The user has to log on the site, then, if this is the first visit, he is asked to evaluate a small number of paintings, then he is free to visit the museum. During his visit, he can get a closer look at some paintings, rate the picture, or continue visiting other corridors. Different corridors are constructed to visit paintings by painter, by epoch, or by color and texture similarity.

The Active Web Museum is our web site, but we are also experimenting our methods on other available ratings databases. In particular, we are studying the automatic classification of users and objects, we have combined collaborative filtering with color and texture analysis of the painting. We have also proposed new evaluation schemes to measure the performance of the prediction in a user-oriented fashion.



### References

- [1] Arnd Kohrs and Bernard Merialdo.  
Improving collaborative filtering with multimedia indexing techniques to create user-adapting web sites.  
In Proceedings of the 7<sup>th</sup> ACM Multimedia Conference, Orlando. ACM, 1999.
- [2] Arnd Kohrs and Bernard Merialdo.  
Using category-based collaborative filtering in the Active WebMuseum.  
In IEEE International Conference on Multimedia and Expo, New York, 2000.

Recommender Systems at the MITRE Corporation  
Frank Linton  
linton@mitre.org

The MITRE Corporation has developed three recommender systems for internal use. In contrast to commercial systems that recommend products to consumers, our systems recommend knowledge to individuals who are focused on performing their job tasks. One system, OWL, recommends new software functionality to individual users based on the functionality the individual's peers have found valuable. Another system, IURS, recommends URLs on the corporate Intranet to users; again based on the URLs the individual's peers have found valuable. The third system, ExpertFinder, locates individuals having some expertise, based on documents they have posted on the corporate Intranet.

In the OWL project (Linton & Schaefer, 2000) we instrumented an application (we instrumented Microsoft Word, but the method would work with any application) and gathered usage data from as many as 50 users at a time over a 3 year period. OWL compared the commands (by *commands* we mean cut, paste, file/open, file/print, etc) used by each individual to those used by the whole group and recommended new learning to each individual based on the difference between the individual's and the group's knowledge. The knowledge of which commands were worth knowing was constructed by the natural workplace activities of the group. It was not acquired by a conventional instructional design process nor transmitted by a human instructor. This instructional approach - monitoring a groups' behavior and generating individualized feedback - is now possible because monitoring is possible in the new networked workplace environment. The concept is not unique, OWL recommends learning in the same manner that commercial websites such as Amazon.com recommend purchases, but it is a radical departure from traditional instructional design, development, and delivery methods.

The MITRE Intranet contains over 800,000 technical and administrative documents. Despite the presence of links, search tools, keywords, and indexes, we suspect information is underutilized, resulting in duplication of effort. For this reason we have developed the Intranet URL Recommender System (IURS), a tool for information consumers. IURS determines which URLs - from a selected topic - one's peers have visited earlier and guides the user to them. If the user identifies him- or herself to IURS, it will report only those URLs s/he has not already visited. Behind the IURS interface is a database of employee demographic and computer ownership data and a large dataset of Intranet server logs that is updated daily. These same data sources are also utilized by two other applications, one, on the Intranet portal, gives employees their current 'Top 10' URL recommendations, another enables Intranet information providers to see the demographics of their viewers.

The ExpertFinder system (Maybury, D'Amore, & House, forthcoming) helps workers find fellow employees who are knowledgeable about a particular subject area (such as "data mining" or "XML"). Individuals make their skills and knowledge known to the ExpertFinder (and therefore to others) by posting work products that utilize their skills and knowledge, typically documents, on the MITRE Intranet. ExpertFinder uses the existing search engine to find documents containing a user-specified keyword or phrase (e.g., "software agents") and uses a commercially available name extraction tool to find authors and names mentioned near the query terms within the documents. Each name is then weighted by several factors and the results are presented to the user in ranked order.

Recommender systems can improve workers' informal learning by collaborative bootstrapping. The main issues we have uncovered in this work are ensuring recommendation quality, developing reliable prototype software, maintaining data quality, and working in a fast-changing environment.

## References

- Linton, F., & Schaefer, H.P. (2000). Recommender Systems for Learning: Building User and Expert Models through Long-Term Observation of Application Use. *User Modeling And User-Adapted Interaction: Special issue on Deployed User/Student Modeling Systems*, 10, (2/3), 181-208.
- Maybury, M., D'Amore, R. and House, D. forthcoming. Automated Discovery and Mapping of Expertise. In M. Ackerman, A. Cohen, V. Pipek, & V. Wulf (eds.). *Beyond Knowledge Management: Sharing Expertise*, Cambridge: MIT Press.

# Recommender Systems Research at NTNU

Our research focus is on using agents to represent mobile customers and services [3] to perform service and product recommendations among a set of peer agents [1], in addition to applying agents as a general approach for targeted advertisements [2].

## Research Institution

[Department of Computer and Information Science](#) (IDI),  
[Norwegian University of Science and Technology](#) (NTNU)  
IDI/NTNU, N-7491 Trondheim, Norway

## People

<b>Associate Professor</b>	Mihhail Matskin
<b>E-Mail</b>	<a href="mailto:misha@idi.ntnu.no">misha@idi.ntnu.no</a>
<b>Homepage URL</b>	<a href="http://www.idi.ntnu.no/~misha/">http://www.idi.ntnu.no/~misha/</a>
<b>Mobile</b>	+47 9 189-7334
Dr. Matskin has works in agent technology, information systems, formal methods for software design and knowledge-based programming. His current research interests include telecommunication-based services, multi-agent system environments and agent-mediated e-commerce.	

<b>Ph.D. Student:</b>	Amund Tveit
<b>E-Mail</b>	<a href="mailto:amundt@idi.ntnu.no">amundt@idi.ntnu.no</a>
<b>Homepage URL</b>	<a href="http://www.idi.ntnu.no/~amundt/">http://www.idi.ntnu.no/~amundt/</a>
<b>Mobile</b>	+47 4 162-6572
Amund Tveit's research interests include agent-based software architecture, mobile commerce, recommender systems/collaborative filtering, information retrieval, web/data mining and mobile interface agents.	

## Research Project

**Name** **EIComAg - The Electronic Commercial Agents Project**  
**URLs** <http://www.elcomag.com/>  
<http://www.jfipa.org/>

## Publications

[1] A. Tveit. "Peer-to-peer based Recommendations for Mobile Commerce". *ACM Mobile Commerce Workshop*, Rome, Italy, July 21, 2001.

[2] M. Matskin. "Collaborative Advertising over Internet with Agents". To appear in *Proceedings of the Intl. Workshop on Web Based Collaboration. WBC 2001*, Munich, Germany, September 3-7, 2001

[3] M. Matskin and A. Tveit. "Mobile Commerce Agents in WAP-Based Services". *Special Issue on Mobile Commerce, Journal of Database Management*, IDEA Group Publishing, July-September 2001, pp 27-35

## Stuart E. Middleton

IAM research group,

Department of Electronics and Computer Science,

University of Southampton, Southampton, S017 1BJ, UK

Email : sem99r@ecs.soton.ac.uk - Web : http://www.iam.ecs.soton.ac.uk/

### Intelligence, agents and multimedia (IAM) group

One of the largest groups of its kind with over 80 researchers, IAM focuses on the design and application of computing systems for complex information and knowledge processing tasks.

IAM is a world leader in the key technologies of agent-based computing, knowledge management, open hypermedia and pervasive computing and their application in the domains of digital libraries and grids.

We see recommender systems as a way to bring together agent, user modelling and machine learning technologies, creating systems grounded in the real world that can provide useful service to real people.

### Past Recommender System Work

#### The MEMOIR project

The MEMOIR system [1] was an open framework supporting a set of agents and hypermedia link services. User web browsing trails allowed MEMOIR to recommend colleagues with similar interests. MEMOIR was installed and evaluated by two end-user organizations.

Subsequent work moved from trails to bookmarks, introducing explicit annotation and replacing the underlying object-oriented database with agents implemented in Prolog.

#### The QuIC project

QuIC [2] explored the idea of dynamically adding multi-destination links to web pages based on the context of the pages. Content-based web page analysis was performed to create a context and this context used to recommend the most similar linkbase. Contextual links were added from the recommended linkbase to the users web pages.

#### The Quickstep project

The Quickstep recommender system [3] explored the acquisition of user profiles by unobtrusive monitoring of browsing behaviour and application of supervised machine-learning techniques.

Specifically, Quickstep addressed the problem of researchers finding relevant on-line research papers from the Internet. A detailed user profile was built from both unobtrusive monitoring of URL browsing and relevance feedback on recommendations provided. Content-based and collaborative filtering techniques were applied to formulate the daily recommendations.

Two identical experiments were conducted in order to evaluate its usefulness as a whole and test two profiling techniques. One profiling technique used a flat list of

topics, the other an “is-a” ontological representation that allowed superclass interest to be inferred from more specific subclasses.

Both empirical studies suggested that the ontological approach was superior. A third, identical trial is planned to allow more rigorous statistical analysis to be applied.

### Current Recommender System Work

#### The Foxtrot project

The Foxtrot recommender system is an evolution of the Quickstep system. It supports a searchable research paper database and offers recommendations based on observed behaviour and relevance feedback.

The major difference is that interest profiles can be visualized, so feedback can be provided directly about a profile. It is expected that a dialogue with the “profiling agent” will improve the quality of the feedback and enhance the users feeling of control over the recommendation process.

An experiment to measure the effectiveness of this enhanced feedback is to be conducted for 1 academic year, using about 150 undergraduates split into two groups. One group will be able to provide relevance feedback and direct profile feedback, the other just relevance feedback. The results will measure the comparative benefit direct feedback provides, and the overall effectiveness of such a recommender system.

### Contact details

Stuart E. Middleton	sem99r@ecs.soton.ac.uk
Wendy Hall	wh@ecs.soton.ac.uk
Dave de Roure	dder@ecs.soton.ac.uk
Nigel R. Shadbolt	nrs@ecs.soton.ac.uk

### REFERENCES

1. De Roure, D. Hall, W. Reich, S. Hill, G. Pikrakis, A. Stairmand, M. MEMOIR – an open framework for enhanced navigation of distributed information, *Information Processing and Management*, 37, 53-74, 2001
2. El-Beltagy, S. Hall, W. De Roure, D. Carr, L. Linking in Context. *The Twelfth ACM Conference on Hypertext and Hypermedia (Hypertext '01)*, 2001, p.151-160
3. Middleton, S.E. De Roure, D. Shadbolt, N.R. Capturing knowledge of user preferences: ontologies in recommender systems, *The First International Conference on Knowledge Capture (K-CAP-2001)*, Victoria, B.C., Canada, Oct 2001

# Multiagent Systems for the Social Competition Among Website Ratings

Matthias Nickles

Department of Informatics, Technical University of Munich  
80290 München, Germany, nickles@cs.tum.edu

[www7.in.tum.de/gruppen/kikog/projects/socionics/construct-s.htm](http://www7.in.tum.de/gruppen/kikog/projects/socionics/construct-s.htm)

## Summary

The self-description of websites and their hyper links referring to other sites usually reflect solely the requirements, opinions and preferences of the respective site owners, there is no explicit relation to the information, service and product offers of competing sites, and the assessment of the site by other users is unknown as well. To provide a “socially enhanced” solution for this problem, our project works towards open multiagent systems which (in contrast to usual filtering-oriented recommender systems) generate multi-dimensional, semantic website ratings which explicitly weigh and combine the possibly conflicting opinions of interacting rating agents which compete in the assertion of their individual ratings against others. These rating agents act as representatives of web surfers, interest communities, other web sites, private or public organizations, and represent their individual ratings in an open discussion forum which is attached to the rated web site. This forum is continuously observed by a rating instance which computes semantically rich, social weighted *general ratings* (of different abstraction levels) from the individual ratings as they show themselves in the agents communications. In contrast to the results of simple majority voting based recommender systems, a general rating makes social structures (like opinions, norms and roles) which emerge from communications *explicit* instead of filtering them regarding a certain user profile. In addition to the presentation of general ratings to the web surfer, these ratings can also be queried by external information agents (like web spiders) and by the rating agents themselves to improve their rating capabilities.

The described approach is based on the technical concept of the so-called *Social System Mirror* (SSM) component, which derives and models the social structures of a multiagent system and makes them visible both to the agents and to the application users. This concept is developed by the *ConStruct* project, which makes use of sociological theories to construct architectures for large adaptive multiagent systems (“artificial societies”) modelled after human society. Take a look at our website for more information and related project publications.

Keywords: Semantic web, multiagent systems, collaborative filtering, socionics

# Recommender Systems Research at NEC Research Institute

**David M. Pennock, Steve Lawrence, Gary William Flake, Brian Whitman**

NEC Research Institute, 4 Independence Way, Princeton, NJ 08540 USA

{dpennock,lawrence,flake,bwhitman}@research.nj.nec.com

<http://www.neci.nec.com/>

At NEC Research Institute, we are investigating new algorithms and applications for recommender systems. We have developed algorithms both for the purely collaborative case, and for combining collaborative and content-based data. We are actively pursuing two major application areas: (1) recommending papers from ResearchIndex ([researchindex.org](http://researchindex.org)), an online digital library of nearly 500,000 computer science papers [1], and (2) recommending music.

Our collaborators include Professor Lyle Ungar, Alexandrin Popescul, and Andrew I. Schein from the University of Pennsylvania; Professor John Riedl, Dan Cosley, and Shyong Lam from the University of Minnesota; Dr. Eric Horvitz from Microsoft Research; Adam Berenzweig and David Birchfield from Columbia University; and Professor C. Lee Giles from Pennsylvania State University.

**Algorithms.** Our *personality diagnosis* algorithm [3] is a pure collaborative filtering algorithm based on a probabilistic model of how people rate items. The algorithm empirically outperforms several previous methods on the EachMovie database of movie ratings, and on user profile data collected from ResearchIndex. The probabilistic framework naturally supports a variety of descriptive measurements—in particular, we consider the applicability of a value of information computation.

In other work [2], we examine some of the theoretical foundations of collaborative filtering algorithms, applying and extending axiomatizations from social choice theory.

Our generalized aspect models [4] offer a unified probabilistic framework for merging collaborative and content-based recommendations. The relative influence of collaboration data versus content data emerges naturally from the given data sources. Experiments on data from ResearchIndex show that appropriate mixture models incorporating secondary data produce significantly better quality recommenders than standard methods.

Recommender systems that rely solely on collaborative data fail when operating from a *cold start*—that is, when recommending items (e.g., first-run movies) that no member of the community has yet seen. In a paper appearing at this workshop [5], we develop several generative probabilistic models that circumvent the cold-start problem by merging multiple forms of content and collaborative data.

We plan to continue to pursue methods for combining content and collaborative data using generative probabilistic models. We are also exploring algorithms for using genre or other clustering information in recommender systems. Additionally,

we are pursuing tools and techniques for scaling and generalizing recommendation algorithms to operate on massive and dynamic data sources.

**Applications.** We are developing a generic framework to allow anyone to evaluate recommendation algorithms using ResearchIndex. Developers can build on a skeleton client which handles communication with ResearchIndex and provides information about user actions and similarity between research papers. We also provide a fully-functional recommender engine using the personality diagnosis algorithm described above. We plan to evaluate recommenders using click-through rates, which provide a number of advantages over standard metrics such as mean absolute error. Standard metrics are often limited in domains where ratings are neither precise nor explicit, as is the case with ResearchIndex. Analyzing click-through rates can also provide an accurate picture of how a recommender impacts user behavior. Click-through is also a more relevant metric for the e-commerce applications typical of recommender systems.

We are also developing a music recommendation system, Minnowmatch, which combines collaborative and content-based recommendations. Minnowmatch uses analysis of the audio content of songs, as well as explicit and implicit user ratings.

## References

- [1] Steve Lawrence, C. Lee Giles, and Kurt Bollacker. Digital libraries and Autonomous Citation Indexing. *IEEE Computer*, 32(6):67–71, 1999.
- [2] D. M. Pennock, E. Horvitz, and C. L. Giles. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *National Conference on Artificial Intelligence*, pages 729–734, 2000.
- [3] D. M. Pennock, E. Horvitz, S. Lawrence, and C. L. Giles. Collaborative filtering by personality diagnosis: A hybrid memory- and model-based approach. In *Conference on Uncertainty in AI*, pages 473–480, 2000.
- [4] A. Popescul, L.H. Ungar, D. M. Pennock, and S. Lawrence. Probabilistic models for unified collaborative and content-based recommendation in sparse-data environments. In *Conference on Uncertainty in AI*, pages 437–444, 2001.
- [5] A. I. Schein, A. Popescul, L.H. Ungar, and D. M. Pennock. Generative models for cold-start recommendations. In *Workshop on Recommender Systems at SIGIR*, 2001.

**Altered Vista**

<http://alteredvista.usu.edu/>

**Researchers**

Mimi Recker (mimi.recker@usu.edu)

Andrew Walker (andy.walker@usu.edu)

Department of Instructional Technology

Utah State University

2830 Old Main Hill

Logan, UT, 84332-2830, U.S.A.

**Description**

Altered Vista is a web based collaborative filtering tool that is geared toward educational web sites (both for students as well as teachers). In its current form, users explicitly submit a detailed rating form for web sites and then request recommendations that are generated using a neighborhood based correlation approach. The system has gone through several user trials and data for the most recent of these is available for download on the project site. Preliminary findings suggest that this particular user population tends to rate with a ceiling effect. While this results in extremely accurate predictions, these predictions do not outperform a simple non-personalized community average.

**Support**

This work reported is partially supported by a grant from the National Science Foundation (NSF NSDL DUE-0085855) and a Utah State University New Faculty Grant.

## **Instructional Architect**

<http://ia.usu.edu/>

### **Researchers**

Mimi Recker (mimi.recker@usu.edu)

Jim Dorward (jimd@coe.usu.edu)

David Wiley (wiley@cc.usu.edu)

Department of Instructional Technology

Utah State University

2830 Old Main Hill

Logan, UT, 84332-2830, U.S.A.

### **Description**

Several initiatives have sparked the creation of digital libraries of on-line learning resources, also called 'learning objects'. While these collections contain some excellent resources, there is little provision for embedding these objects into new curriculum. Instructional Architect (IA) is an NSF funded project in development at Utah State University, which supplies a service layer on top of digital libraries. IA currently allows educational practitioners to search for learning objects and combine them into finished instruction. Work is underway to use collaborative filtering in recommending learning objects to users of the system based on the assumption that use of a learning object is an implicit vote of preference for that specific resource.

### **Support**

This work reported is partially supported by a grant from the National Science Foundation (NSF NSDL DUE-0085855).

**Phil Rennert**

EC Wise Machine Learning Center  
408 Saybrooke View Dr.,  
Gaithersburg MD 20877  
301-527-1230  
prennert@ecwise.com

We do Web usability solutions, and we use a number of machine learning approaches, including collaborative filtering, to predict elements of page and site design which will be most useful and usable to a target audience. Predictions are based initially on demographics, later on user feedback and survey results; so we're yet another group grappling with the problem of how best to mix demographics and recommendations.

In a previous job, I wrote a collaborative filtering system to recommend colleges to high school students, based on the list of other colleges where they'd chosen to apply, for a Web-based college application service. From holdout test set results, the filter outperformed demographic clustering predictions and unpersonalized recommendations.

## GroupLens Research: Recommender System Interfaces and Algorithms

*PIs: John Riedl and Joseph A. Konstan*

*E-mail: {riedl, konstan}@cs.umn.edu*

*<http://www.grouplens.org>*

*4-192 EE/CS Building, University of Minnesota, Minneapolis, MN-55455*

The GroupLens Research group has been working on Recommender Systems since we began exploring collaborative filtering in 1992. Our work has focused on two interlocking questions: what are the best interfaces for presenting recommender systems to users, and what are the best algorithms for producing recommendations. In this work we have run two major research systems: the original GroupLens for Usenet news (currently on hiatus), and MovieLens, a recommender for movies ([www.movielens.org](http://www.movielens.org)) with over 95,000 users and 6.7 million ratings.

In our research on interfaces we have recently been exploring new types of input from users that enable new ways to use recommenders, such as group recommenders, explanation systems, email-based recommenders, and ephemeral interest recommenders. We will briefly discuss the first two of these here; please write us to learn about the others.

Group recommenders make recommendations for items that are best for a whole group of people, though they may not be best for any one of the individuals. For instance, the group recommender for MovieLens can help a group of people find a movie to go to. The group recommendation is produced by combining individual recommendations with the goal of maximizing the “social good” for the group.

Explanation systems for recommenders provide an interface for users to understand better how it is that the recommender works, and why the recommender is suggesting a certain item. In MovieLens, our most common question is something along the lines of “Your recommendations are almost always exactly right. MovieLens really understands me. Therefore, how can you possibly think I would like *The Shining*?” The explanation can point out that people the user has agreed with about other movies like *The Shining*, can list other movies that are rated similarly to *The Shining*, or can otherwise explain what MovieLens is “thinking”.

In our research on algorithms we have been exploring dimension reduction techniques, including singular value decomposition, we have studied the sensitivity of nearest neighbor recommenders to the most important parameters that influence their behavior, we have studied the metrics that are used to measure recommenders, and we have explored item-based recommenders. We will briefly discuss the first of those here.

Dimension reduction techniques have potential for recommenders both because they reduce the search space for finding good neighbors, or good items to recommend, and because they may help uncover latent relationships between the users or items that might be hard to discover in the high-dimensional space of the recommenders. Our research on singular value decomposition (SVD) has demonstrated that for some recommender tasks it offers much faster online performance, while producing results that are as good as the best published results from other algorithms.

We are proud of our former Ph.D. students who are now actively pursuing recommender systems research in universities: Mark Claypool at Worcester Polytechnic Institute, Jon Herlocker at Oregon State University, and Ben Schafer at the University of Northern Iowa.

Dr. J. Ben Schafer  
Department of Computer Science  
University of Northern Iowa  
Cedar Falls, IA 50614  
Ben.Schafer@uni.edu

### **Prior Work**

As a member of the GroupLens Research Project at the University of Minnesota I conducted research focusing on three aspects of recommender systems:

- ◆ A Taxonomy for Recommender Systems in Electronic Commerce

An analysis of how recommender systems have been applied in the field of electronic commerce was conducted. From this analysis, a taxonomy for these systems was built, classifying them based on the functional I/O, the recommendation method, and other design issues such as degree of personalization and delivery method. Additionally, we identified six application models for recommender systems in electronic commerce. Finally, we conducted an extensive survey of electronic commerce sites and considered patterns related to which application models exist in which product domains.

- ◆ Filterbots

Information filtering agents (IF) and collaborative filtering (CF) both attempt to alleviate information overload by identifying which items a user will find worthwhile. Each of these techniques has advantages and limitations. This research demonstrated that the two can be beneficially combined. A CF framework was used to combine personal IF agents and the opinions of a community of users to produce better recommendations than either agents or user can produce alone. Additionally, it was concluded that a personal CF system consisting solely of a user's personal IF agents produced more accurate recommendations than individual agents alone or the recommendations produced by other combination techniques.

- ◆ Meta-recommender Systems

We have introduced a new class of recommender systems we classify as *meta-recommender systems*. These systems provide users with personalized control over the generation of a single recommendation list formed from the combination of rich recommendation data from multiple information sources and recommendation techniques.

As part of this research, we have presented the design and implementation of a meta-recommender framework named the *MetaLens Recommendation Framework* (MLRF). MLRF is based on an evaluation system model used by an increasing number of Internet sites to help consumers find items of interest. These sites allow consumers to identify a domain of purchase (i.e., a digital camera, computer, or automobile) and narrow the list of products within that domain. Consumers may indicate the features that they would be interested in, and the importance of these features in their final decision. Sites turn these requirements into search queries, use information filtering techniques over the attributes of known products in the category, and return ranked lists of "recommended" products. MLRF extends this model by evaluating how well individual consumers will *like* the items and including this evaluation in the recommendation process.

Research surrounding MLRF has focused on user preferences in meta-recommendation format, user perceived value of meta-recommenders, and usage analysis of three publicly available meta-recommenders built within MLRF for the domain of movies. Analysis of public use of these systems demonstrates that users like, and often prefer, these systems to more "traditional" recommenders. While acceptance comes at a slow pace, users who customized a system were more likely to return to use the system again. Finally, while the quantity and type of recommendation data preferred varies widely from user to user, analysis demonstrates that users want access to as much recommendation data as possible.

### **Current/Future Work**

Having recently graduated from the University of Minnesota, I am currently in the process of establishing my research lab at the University of Northern Iowa. Research plans for the immediate future focus on interface and algorithmic issue for meta-recommenders. These include the creation of modular meta-recommenders, user-defined sources for recommendation data, interfaces for meta-recommendations through direct manipulation, and metrics for comparing the "accuracy" of different algorithms for meta-recommendations.