

AI Techniques for Personalized Recommendation

Proposal for a full-day AAAI 2004 tutorial

John Riedl and **Anthony Jameson** and **Joseph Konstan**

1 Brief Description of the Tutorial

Personalized recommendation of products, documents, and collaborators has become an important way of meeting user needs in commerce, information provision, and community services, whether on the web, through mobile interfaces, or through traditional desktop interfaces. This tutorial first reviews the types of personalized recommendation that are being used commercially and in research systems. It then systematically presents and compares the underlying AI techniques, including recent variants and extensions of collaborative filtering, demographic and case-based approaches, and decision-theoretic methods. The properties of the various techniques are compared within a general framework, so that participants learn how to match recommendation techniques to applications and how to combine complementary techniques.

The full-day format makes it possible to include a session in which participants actively work together on a concrete problem, as well as in-depth discussion of application contexts, case studies, and key social issues.

The tutorial presupposes a general knowledge of AI. Some previous familiarity with issues of personalized recommendation is desirable but not essential.

2 Detailed Outline

This tutorial will be an updated version of a successful full-day tutorial that was held at AAAI 2002 in Edmonton. We also gave a half-day version of the tutorial at IJCAI 2003 in Acapulco, since the IJCAI tutorial program could not accommodate full-day tutorials. Although the half-day version was also well received, we very much want to return to the full-day format for AAAI 2004, for reasons to be detailed below.

Table 1 shows the proposed schedule for the tutorial. The following subsections comment on the various parts of the tutorial shown in the table.

2.1 Part 1: Introduction, Applications, and Interfaces

Introduction

- Basic concepts and motivation
- Tutorial goals and preview
- Brief history of personalized recommendation
- Background of presenters and participants

Recommender Application Space

- Overview of several dimensions along which recommender applications vary (e.g., purpose of recommendations)
- Initial discussion of privacy issues

Applications and Interfaces

- Discussion of a representative set of types of application and interface from three areas:
 - e-commerce
 - document recommendation
 - off-web interfaces

Case Studies of Applications and Interfaces

- Several in-depth discussions of particular applications and interfaces, in the form of:
 - live demonstrations
 - screen shots of important systems that are no longer available

In addition to illustrating the concepts introduced in the previous section, the case studies raise further specific issues and encourage interactive discussion.

Overview of the Major Approaches to Personalized Recommendation

- Initial presentation of a general integrative framework for all of the techniques to be discussed in this and the following sections

The five major types of recommendation technique are distinguished on an abstract level, and a preview is given of the ways in which they complement each other by having different sets of typical strengths and limitations. This fact in turn points to the value of hybrid methods that combine the strengths of different approaches.

2.2 Part 2: Techniques Based on Collaborative Filtering:

It is explained that collaborative filtering will be discussed in more detail than each of the other approaches, because many of the general issues are introduced for the first time in this section and because more experience has been acquired in this area with the practical issues that arise with deployed systems.

Table 1. Proposed tutorial schedule.

(This schedule presupposes that the sessions and coffee breaks begin at the same times of day as at AAAI 2002.)

Morning

9:00–9:20	Introduction
9:20–9:35	Recommender Application Space
9:35–10:15	Applications and Interfaces
10:15–10:45	Case Studies of Applications and Interfaces
<i>Coffee Break</i>	
11:15–11:30	Overview of Recommendation Techniques
11:30–12:30	Collaborative Filtering Techniques
12:30–1:00	Collaborative Filtering Case Studies

Afternoon

<i>Lunch break</i>	
2:00–2:45	Non-CF Techniques
2:45–3:15	Non CF Case Studies
3:15–3:45	Hybrid Techniques
<i>Coffee break</i>	
4:15–5:00	Designing Recommender Applications
5:00–5:45	Advanced Topics
5:45–6:00	Conclusions, Questions

Collaborative Filtering Techniques

- Basic K-nearest-neighbor algorithms
- Item-item collaborative filtering
- Dimensionality reduction algorithms
- Explaining recommendations based on collaborative filtering

Collaborative Filtering Case Studies

- The case studies serve a function similar to that of the case studies for applications and interfaces, but the emphasis is on the recommendation techniques.

2.3 Part 3: Techniques Not Based Solely on Collaborative Filtering

Non-CF Recommendation Techniques

- Content-based methods, including those that make use of case-based reasoning or text classification methods
- Demographically based methods, which base recommendations on personal characteristics of the customer
- Utility-based methods, which predict the value of items to each individual user on the basis of a model of the user's preferences, which may have been elicited explicitly
- Knowledge-based methods, which incorporate some sort of domain knowledge to match users' requirements with the properties of items

Case Studies of Non-CF Recommendation Techniques

- In-depth discussions of particular systems that embody some of the techniques introduced in the previous section

Hybrid Recommendation Techniques

- General overview of types of hybrid techniques
- Examples of applications of hybrid techniques

2.4 Part 4: Applying and Extending What Has Been Learned So Far

Designing Recommender Applications

- Instructions on how to choose a particular application and work out several high-level aspects of its design
- Discussion in groups of about 5 participants, with occasional advice from the presenters
- Brief reports from the groups on their ideas and the problems they encountered

Advanced Topics

- Presentation and discussion of selected recent developments in personalized recommendation, such as:
 - special issues raised by recommender systems for groups, such as the issue of nonmanipulable preference aggregation
 - extended uses of data obtained from recommender systems
 - issues in mobile/wireless recommender systems
 - meta-level recommender systems
 - anthropomorphism in recommender systems
 - agent participation in recommender systems

Conclusions, Questions

- Discussion of any further issues raised by the participants

3 Necessary Background and Potential Target Audience

While the tutorial presupposes a general knowledge of AI, it does not presuppose any particular knowledge of personalized recommendation. Participants at the AAAI 2002 and IJCAI 2003 tutorials ranged from persons seeking an introductory overview to well-known researchers with important publications in the field of personalized recommendation. All of these participants can benefit because of the integrative nature of the tutorial, which serves the objective "Present a

novel synthesis combining distinct lines of AI work". This perspective results in the inclusion of a much broader range of material than any single participant is likely to be familiar with.

Whereas previous tutorials have usually focused on one particular type of learning or inference technique, this tutorial covers techniques ranging from all major extensions of collaborative filtering to case-based reasoning, decision-theoretic approaches, and various types of hybrid model. These techniques are discussed in relation to the wide variety of research and commercial applications that employ AI-based recommendation. The algorithms, applications, and interfaces presented are illustrated through a set of case studies drawn from both research and commercial systems.

The target audience comprises (a) attendees from industry who seek an authoritative update on recommendation technology; (b) AI researchers who see personalized recommendation as an attractive application area; and (c) researchers already working on one type of recommendation technology who would like to learn more about alternative approaches. After taking this tutorial, attendees will (1) be familiar with a broad set of recommendation techniques and be able to identify appropriate techniques for specific recommendation challenges; (2) be familiar with the major recommender interfaces and able to match appropriate interfaces to specific user and designer goals; and (3) be familiar with a large set of recommendation applications, and able to use these as models when designing and developing future applications.

4 Full-Day Format

It is in principle possible to present much of this material in a lecture format within a half-day tutorial, as we did at IJ-CAI 2003. But if this is done, some of the most distinctive features of the tutorial, which constituted the original motivation for developing it, are lost:

1. The presentation and discussion of case studies from the world of deployed systems, which illustrate how personalized recommendation involves many more AI challenges than just that of generating accurate recommendations.
2. Lively interactive discussion with the participants of issues of particular interest to them, as they arise throughout the tutorial.
3. The group exercise, in which participants get a chance to apply some of what they have learned while it is still fresh in their minds.

The full-day tutorial at AAAI 2002 demonstrated that these distinctive features really are realized within this format and that they are highly appreciated by the participants. It is therefore questionable whether it would be worthwhile for the three presenters to give the severely reduced half-day version.

5 Interest for a AAAI Audience

Personalized recommendation has been a lively application area for AI techniques for several years now. For example, every IJCAI and AAAI conference includes a number

of papers and posters from this broad area. But an unnecessary degree of fragmentation into subcommunities can be observed, each of which focuses on one type of AI technique, such as: classical collaborative filtering, case-based reasoning, decision-theoretic methods, or techniques for preference elicitation. Moreover, many researchers in this area have little experience with the practical issues that arise with the use of personalized recommendation – issues which place some strong constraints on the use of AI techniques.

This tutorial was developed with the explicit goal of overcoming this fragmentation. The very positive feedback from the participants at AAAI 2002 indicates that the goal was achieved.

6 Resumes of the Presenters

One of the strengths of this tutorial, particularly for the AAAI audience, is the fact that it brings together broad experience with AI techniques for user-adaptive systems (Jameson) and extensive research and industrial experience in the specific personalization subarea of recommender systems based on collaborative filtering (Konstan and Riedl). The result is a tutorial with extensive coverage of recommendation techniques supported by a rich collection of application examples and design experience.

6.1 John Riedl and Joe Konstan

Addresses

John Riedl, Associate Professor
Department of Computer Science and Engineering
University of Minnesota
5-211 EE/CS Building
200 Union Street SE
Minneapolis, MN 55455
riedl@cs.umn.edu

<http://www-users.cs.umn.edu/~riedl/>

Phone: (+1)(612) 624-7372

Joseph A. Konstan, Associate Professor
Department of Computer Science and Engineering
University of Minnesota
4-192 EE/CS Building
200 Union Street SE
Minneapolis, MN 55455
konstan@cs.umn.edu

<http://www-users.cs.umn.edu/~konstan/>

Phone: (+1)(612) 625-1831

Background in the Tutorial Area

John Riedl co-founded the GroupLens recommender system project with Paul Resnick in 1992, and he and Joe Konstan have been codirecting the project since 1995. GroupLens was one of the earliest recommender systems projects, and is now one of the most widely known. Riedl and Konstan also co-founded Net Perceptions, a company that has commercialized the results of their research since 1996. They have published broadly in the area of recommender systems, as is indicated by the following selection of publications:

- Cosley, D., Lam, S. K., Albert, I., Konstan, J., & Riedl, J. (2003). Is seeing believing? how recommender systems influence users' opinions. In L. Terveen, D. Wixon, E. Comstock, & A. Sasse (Eds.), *Human factors in computing systems: CHI 2003 conference proceedings* (pp. 585–592). New York: ACM.
- Herlocker, J., Konstan, J., Borchers, A., & Riedl, J. (1999). An algorithmic framework for performing collaborative filtering. *Proceedings of the 1999 Conference on Research and Development in Information Retrieval*.
- Herlocker, J., Konstan, J., & Riedl, J. (2000). Explaining collaborative filtering recommendations. *Proceedings of the 2000 Conference on Computer-Supported Cooperative Work*, Philadelphia, PA, pp. 241–250.
- Konstan, J., Miller, B., Maltz, D., Herlocker, J., Gordon, L., & Riedl, J. (1997). GroupLens: Applying collaborative filtering to Usenet news. *Communications of the ACM*, 40(3), 77–87.
- Miller, B., Riedl, J., & Konstan, J. (1997). Experiences with GroupLens: Making Usenet useful again. *Proceedings of the 1997 Usenix Winter Technical Conference*.
- Resnick, P., Iacovou, N., Sushak, M., Bergstrom, P., & Riedl, J. (1994). An open architecture for collaborative filtering of netnews. *Proceedings of the 1994 Conference on Computer-Supported Cooperative Work*.
- Sarwar, B., Karypis, G., Konstan, J., & Riedl, J. (2000). Analysis of recommender algorithms for e-commerce. *Proceedings of the ACM E-Commerce 2000 Conference*.
- Sarwar, B., Konstan, J., & Riedl, J. (2001). Distributed recommender systems: New opportunities for internet commerce. In S. M. Rahman & R. Bignall (Eds.), *Internet commerce and software agents: Cases, technologies and opportunities*. Hershey, PA: Idea Group Publishing.
- Schafer, J. B., Konstan, J., & Riedl, J. (2001). E-commerce recommendation applications. *Journal of Data Mining and Knowledge Discovery*, 5(1/2), 115–152.
- Schafer, J. B., Konstan, J., & Riedl, J. (1999). Recommender systems in e-commerce. *Proceedings of the ACM Conference on Electronic Commerce*.

They also recently published the following book:

- Riedl, J., Konstan, J., & Vrooman, E. (2002). *Word of mouse: The marketing power of collaborative filtering*. New York: Warner Books.

In addition to the AAAI 2002 and IJCAI 2003 tutorials on which the presently proposed tutorial is based, Joe Konstan and John Riedl have taught several versions of a tutorial on recommender systems, at the 1996, 2000, and 2002 ACM conferences on Computer-Supported Cooperative Work (CSCW), at the 2000 ACM E-Commerce conference, and at the 2003 ACM conferences on Intelligent User Interfaces (IUI) and Computer-Human Interaction (CHI). The tutorial has evolved with each offering, but the most recent version includes an introduction to a variety of algorithms and a set of ten case studies of systems in use. The CSCW 2000

tutorial received the top overall score of tutorials at the conference.

Evidence of Teaching Experience

Joe Konstan has been on the faculty of the Computer Science and Engineering Department of the University of Minnesota since 1992. He has taught many University Computer Science and Software Engineering courses in the areas of human-computer interaction, programming, and computer graphics. He has consistently earned among the top student evaluations in the department, and was awarded the George Taylor award as the best young teacher in the college in 1999. He has also taught many industry short courses.

John Riedl has been on the faculty of the Computer Science and Engineering Department of the University of Minnesota since 1990. He has taught over twenty University courses in the areas of object-oriented programming, operating systems, and related areas. He has earned the award as the department's best teacher three times, and was awarded the George Taylor award as best young teacher in the college in 1996.

Evidence of Scholarship in AI or Computer Science

Konstan and Riedl are both tenured associate professors in the Department of Computer Science and Engineering at the University of Minnesota. In addition to their work together on recommender systems listed above, Konstan has published extensively in human-computer interaction and multimedia, and Riedl has published extensively in collaborative systems, scientific databases, and visualization. Three example articles from each, apart from the collaborative filtering articles listed above, include:

Joseph Konstan:

- Carlis, J., & Konstan, J. (1998). Interactive visualization of serial periodic data. *Proceedings of UIST98*, San Francisco, pp. 29–38.
- Gustafson, T., Schafer, J., & Konstan, J. (1998). Agents in their midst: Evaluating user adaptation to agent-assisted interfaces. In J. Marks (Ed.), *IUI98: International Conference on Intelligent User Interfaces* (pp. 163–170). New York: ACM.
- Konstan, J., Lichtblau, L., Sturm, P., & McLeod, J. (1997). ISAP: A web-based pharmacology course and study environment. *Computers and Education*.

John Riedl:

- Mashayekhi, V., Drake, J., Tsai, W., & Riedl, J. (1993). Distributed, collaborative software inspection. *IEEE Software*, 10(5).
- Claypool, M., Riedl, J., Carlis, J., Wilcox, G., Elde, R., Retzel, E., Georgeopolous, A., Pardo, J., Ugurbil, K., Miller, B., & Honda, C. (1995). Network requirements for 3D flying in a zoomable brain database. *IEEE JSAC*, 13(5), 816–827.

Chi, E., Riedl, J., Barry, P., & Konstan, J. (1998). Principles for information visualization spreadsheets. *IEEE Computer Graphics and Applications*.

6.2 Anthony Jameson

Addresses

Anthony Jameson, Principal Researcher / Adjunct Professor
German Research Institute for Artificial Intelligence (DFKI)
/ International University in Germany
Stuhlsatzenhausweg 3
66123 Saarbrücken, Germany
jameson@dfki.de
<http://dfki.de/~jameson/>
Phone: (+49)(681) 302-5079
Fax: (+49)(681) 302-5020

Background in the Tutorial Area

Jameson's first publication on what is now called *personalized recommendation* appeared in the proceedings of IJCAI 1983. Since then he has published widely on personalization and user-adaptive systems more generally, as is illustrated by the following sample of publications:

- Jameson, A. (1983). Impression monitoring in evaluation-oriented dialog: The role of the listener's assumed expectations and values in the generation of informative statements. *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, Karlsruhe, pp. 616–620.
- Jameson, A. (1989). But what will the listener think? Belief ascription and image maintenance in dialog. In A. Kobsa & W. Wahlster (Eds.), *User models in dialog systems* (pp. 255–312). Berlin: Springer.
- Jameson, A., Schäfer, R., Simons, J., & Weis, T. (1995). Adaptive provision of evaluation-oriented information: Tasks and techniques. In C. Mellish (Ed.), *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence* (pp. 1886–1893). San Mateo, CA: Morgan Kaufmann.
- Jameson, A. (1996). Numerical uncertainty management in user and student modeling: An overview of systems and issues. *User Modeling and User-Adapted Interaction*, 5, 193–251.
- Jameson, A. (2003). Adaptive interfaces and agents. In J. Jacko & A. Sears (Eds.), *Human-computer interaction handbook* (pp. 305–330). Mahwah, NJ: Erlbaum.

The last two publications are broad survey articles.

A complete list of Jameson's publications since 1993, most of which are available electronically, can be accessed via his web homepage.

In addition to the AAAI 2002 and IJCAI 2003 tutorials on which the presently proposed tutorial is based, Jameson has recently presented the following tutorials that are closely related to the topic of the proposed tutorial:

- *Personalization for E-Commerce*. UM 2001, Eighth International Conference on User Modeling, Sonthofen, Germany. (Slides available from <http://dfki.de/~jameson/ec-ttrl/ec-ttrl.pdf>.)
- *Systems That Adapt to Their Users*. IJCAI 2001, Seventeenth International Joint Conference on Artificial Intelligence, Seattle. (Slides available from <http://dfki.de/~jameson/ttrl-notes-ijcai01/>.)
- *Designing Systems That Adapt to Their Users* (full-day tutorial). CHI 2001, Conference on Human Factors in Computing Systems, Seattle; CHI 2002, Minneapolis; UM 2003, Ninth International Conference on User Modeling, Johnstown.
- *Designing User-Adaptive Systems*. IUI 2001, International Conference on Intelligent User Interfaces, Santa Fe, New Mexico.
- *User-Adaptive Systems: An Integrative Overview*. UM 1999, Seventh International Conference on User Modeling, Banff, Canada and IJCAI 1999, Sixteenth International Joint Conference on Artificial Intelligence, Stockholm.

Evidence of Teaching Experience

Jameson has taught three university courses on personalization: one in the summer of 2001 at the International University in Germany (summer of 2001) and two at Saarland University (fall/winter 1998/1999, and fall/winter 2000/2001). More generally, he has taught university courses since 1986 on various topics in AI, cognitive science, and human-computer interaction at the University of Nijmegen, Saarland University, and the International University in Germany. Though affiliated primarily with DFKI, in 2001 he was appointed Adjunct Associate Professor at the International University in Germany in recognition of teaching excellence.

Evidence of Scholarship in AI or Computer Science

Most of Jameson's papers cited above include a substantial AI component. Three other recent papers presented at AI conferences are the following:

- Jameson, A., Hackl, C., & Kleinbauer, T. (2003). Evaluation of automatically designed mechanisms. *Proceedings of the First Bayesian Modeling Applications Workshop at the Nineteenth Conference on Uncertainty in Artificial Intelligence*, Acapulco, Mexico.
- Jameson, A., & Wittig, F. (2001). Leveraging data about users in general in the learning of individual user models. In B. Nebel (Ed.), *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence* (pp. 1185–1192). San Francisco, CA: Morgan Kaufmann.
- Wittig, F., & Jameson, A. (2000). Exploiting qualitative knowledge in the learning of conditional probabilities of Bayesian networks. In C. Boutilier & M. Goldszmidt (Eds.), *Uncertainty in Artificial Intelligence: Proceedings of the Sixteenth Conference* (pp. 644–652). San Francisco: Morgan Kaufmann.

The following three papers, each of which has an AI component, were awarded best-paper prizes at their respective

conferences:

- Müller, C., Großmann-Hutter, B., Jameson, A., Rummer, R., & Wittig, F. (2001). Recognizing time pressure and cognitive load on the basis of speech: An experimental study. In M. Bauer, P. Gmytrasiewicz, & J. Vassileva (Eds.), *UM2001, User Modeling: Proceedings of the Eighth International Conference* (pp. 24–33). Berlin: Springer.
- Berthold, A., & Jameson, A. (1999). Interpreting symptoms of cognitive load in speech input. In J. Kay (Ed.), *UM99, User modeling: Proceedings of the Seventh International Conference* (pp. 235–244). Vienna: Springer Wien New York.
- Jameson, A., Schäfer, R., Weis, T., Berthold, A., & Weyrath, T. (1999). Making systems sensitive to the user's time and working memory constraints. In M. Maybury (Ed.), *IUI99: International Conference on Intelligent User Interfaces* (pp. 79–86). New York: ACM.

7 Example Slides

The following four pages show representative samples of slides from several parts of the IJCAI 2003 tutorial. The complete slides were made available to the participants in electronic form via the unlinked URL <http://dfki.de/~jameson/pdf/ijcai03-tutorial.pdf>, where they are still available.

Introduction and Application Space Section Contents

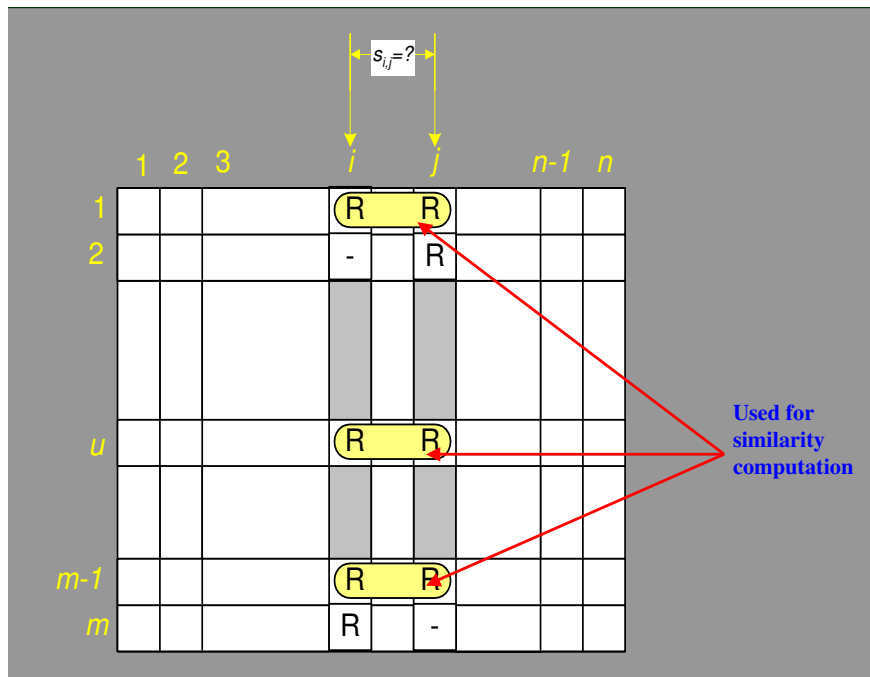
3

<i>What Are Recommender Systems?</i>	5
Preview	5
The Problem: Overload	6
Too much stuff!	7
Recommenders	8
Scope of Recommenders	9
Wide Range of Algorithms	10
<i>Tutorial Goals and Outline</i>	11
Goals	11
Outline	12
<i>History of Recommender Systems</i>	13
The Early Years	13
Information Retrieval	14
Information Filtering	15
Collaborative Filtering	16
Automated CF	17
Usenet Interface	18
ACF Blossomed	19
Today	20
<i>Introductions</i>	21
Presenters	21
About You	22
<i>Recommender Application Space</i>	23
Recommender Application Space	23
Domains of Recommendation	24
Google: Content Example	25
half.com	26
Purposes of Recommendation	27
Buy.com Customers Also Bought	28
OWL Tips	29
Epinions Sienna Overview	30
ReferralWeb	31
Whose Opinion?	32
Wine.com Expert Recommendations	33
Priceline.com Screenshot	34
PHOAKS	35
Personalization Level	36
Lands' End	37
Brooks Brothers	38
CDNOW Album Advisor	39
CDNOW Album Advisor Recommendations	40
My CDMOW	41
Privacy and Trustworthiness	42
Interfaces	43

Figure 1. Section table of contents for the first section.

87

Item Similarities



Item-Item Matrix Formulation

88





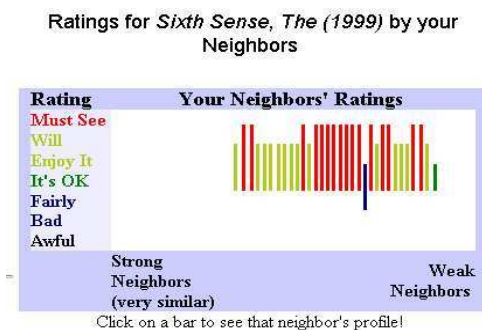
Figure 2. Two of several slides illustrating the method of item-item collaborative filtering.

Most Compelling Interfaces



- *Simple* visual representations of neighbors ratings
- Statement of strong previous performance
"MovieLens has predicted correctly 80% of the time for you"

Less Compelling Interfaces



- Anything with even minimal complexity
 - More than two dimensions
- Any use of statistical terminology
 - Correlation, variance, etc.

Figure 3. Slides presenting results of a study on the explanation of collaborative filtering recommendations.

Formalization (2)

Representation of preferences

- \mathcal{U} 's preferences: soft constraints on the values of attributes
- Constraint function:
 - $C_i(v) : \text{dom}(A_i) \rightarrow [0, 1]$
- 0 = "fully satisfied", 1 = "fully unsatisfied"

Simplifying assumption here

- Preferences are *additive independent*:
 - A preference concerning one attribute does not depend on the level of any other attribute

Formalization (3)




Figure 4. Part of a case study of utility-based recommendation.