

Beyond Idiot Savants: Recommendations and Common Sense

Michael J. Pazzani
Department of Informatics
School of Information and Computer Science
University of California, Irvine
Irvine CA 92612
pazzani@ics.uci.edu

ABSTRACT

The current generation of recommendation systems exhibits little if any common sense. While adept at finding patterns in purchase data, such systems are plateauing well below the goal of having intelligence agents be analogous to human concierges.

Keywords

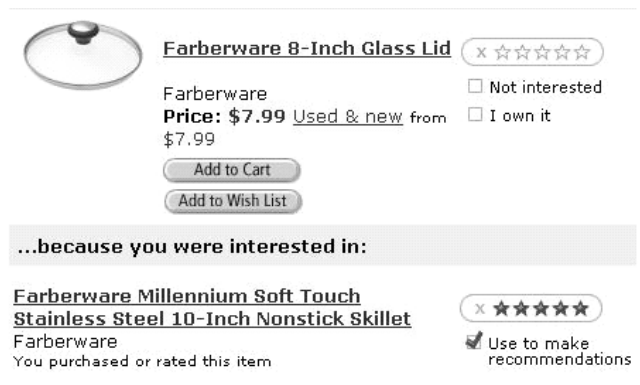
Recommendation systems, personalization, intelligence.

INTRODUCTION

In an early paper on intelligent agents, Etzioni and Weld [1] argued that intelligent agents should be like concierges. However, today's deployed recommendation systems fall short of this goal and are more like idiot savants in that they rely on one or a small number of exceedingly narrow algorithms.

THE PROBLEM

Amazon.com recently recommended that I purchase an 8-inch lid because I had purchased a 10-inch skillet as illustrated in Figure 1. A system that recommends an 8-inch lid for a 10-inch skillet is missing the common sense of sales clerk on his or her first day on job. The problem is not restricted to e-commerce sites, for example the San Diego Union Tribune wireless news site has deployed a content-based recommendation system [2] that recently recommended that I read this story: "Lowe will start Game 7 for Red Sox" because I had recently read a similar story: "Lowe resurrects career with Game 7 gem." Although the two stories have many words in common (i.e. have very similar term vectors), a system that that would recommend an earlier story because one had read the later story is clearly not exhibiting common sense.



Farberware 8-Inch Glass Lid x ☆☆☆☆☆
Farberware Not interested
Price: \$7.99 Used & new from I own it
\$7.99
Add to Cart
Add to Wish List

...because you were interested in:

Farberware Millennium Soft Touch Stainless Steel 10-Inch Nonstick Skillet x ☆☆☆☆☆
Farberware Use to make recommendations
You purchased or rated this item

Figure 1. Recommending a 8-inch lid for a 10-inch skillet.

Of course, one could construct a scenario in which these recommendations do make sense. For example, if one had a complete set of Faberware and damaged a skillet and a lid while moving, it would make perfect sense to suggest purchasing two items of different sizes. However, today's recommendation systems don't even have the common sense to construct such scenarios.

HYBRID ALGORITHMS AREN'T THE SOLUTION

Many have recognized the shortcomings of individual recommendation algorithms (e.g. [3-5]) and have advocated hybrid algorithms combining two or more algorithms in the hopes that the strengths of one complement the weaknesses of the other. However, none of the hybrid systems have attempted to encode common sense knowledge, i.e., today's recommendation systems don't know what they are doing [6].

The problem is that the recommendation systems do not have the common sense to recognizing the user's goals and relating those goals to the recommendation nor can they explain how the recommendations would help satisfy the user's goals. A good concierge would have these abilities.

Other data mining algorithms share the failure of today's recommendation system. For example, through detailed analysis of sales data a system used by a department store might learn that ski jackets sell better in Colorado than

Florida. However, a system that lacks the common sense to explain why this pattern holds would be of little use when the first store opens in Utah or Hawaii.

GOALS AND COMMON SENSE

Today's recommendation systems, although useful do not even approximate the utility of a concierge. To achieve the next level of intelligence, recommendations systems will need at least three capabilities missing from most of today's deployed systems:

- Understanding the user's goals, whether stated by the user or inferred from the user's behavior.
- Representing common sense knowledge that indicates how various actions, such as purchasing items or obtaining information, relate to these goals.
- Integrating knowledge and data from disparate sources. Today's systems are almost always deployed within a single web site while the web provides a vast network of independently developed information resources.

An important benefit of embodying recommendation systems with explicit representation of goals is that recommendations agents would be able to interoperate across sites (cf. [7]). A study of Internet usage revealed that users visit on average 10 sites per session (see <http://www.it-analysis.com/article.php?articleid=1660>) to achieve their goals. An agent that understood the content of such sites would prove valuable in making recommendations by synthesizing information from multiple sources. To give one example, I recently had business in Australia and decided to take a vacation as well. After exploring options, I decided to plan the vacation after my business rather than before because that would allow me to observe a meteor shower on a moonless night from an island with little light pollution. While I had to search manually to construct this plan, I'd expect an intelligent agent that understood my preferences by analyzing my previous vacations to emulate this decision.

Some initial work along the lines advocated here has shown promise. Kim's work on InfoQuilt [8] has demonstrated that ontologies and semantic web allow for personalization across different information sources. Babaian [9] has shown how a declarative representation of preconditions and effects of a system's actions enhances a personalization system. Lieberman et al. [10] have demonstrated the utility of incorporating common-sense reasoning into a variety of applications ranging from organizing digital photos to personalizing the selection of music.

CONCLUSION

Today's recommendation systems operate for the most part by detection correlations between the activity of different users or among the features that describe objects a user likes. However, without explicit representation of user's goals and an explicit representation of world knowledge, such systems lack the ability of a concierge who can explain why there may be correlations in the data and generalize these explanations to new situations.

REFERENCES

1. Etzioni, O & Weld, D. (1995). *Intelligent Agents on the Internet: Fact, Fiction, and Forecast*. IEEE Expert 10(3): 44-49.
2. Billsus, D., Brunk, C., Evans, C., Gladish, B. & Pazzani, M. (2002). *Adaptive Interfaces for Ubiquitous Web Access*, CACM p 34-38.
3. Burke, R. (2002). *Hybrid Recommender Systems: Survey and Experiments*. User Modeling and User-Adapted Interaction v.12 n.4, p.331-370, November 2002
4. Melville, P. Mooney, R. and Nagarajan, R. (2001). *Content Boosted Collaborative Filtering*. Proceedings of the SIGIR-2001 Workshop on Recommender Systems, New Orleans, LA, September 2001
5. Pazzani, M. (1999). *A Framework for Collaborative, Content-Based and Demographic Filtering*. Artificial Intelligence Review. 13(5-6): 393-408
6. Brachman, Ronald J. (2002). *Systems that Know What They're Doing*, IEEE Intelligent Systems (November/December): 67-71.
7. Kuno, H. & Sahai, A. (2003). *My Agent Wants to Talk to Your Service: Personalizing Web Services through Agents*. In B. Burg, J. Dale, T. Finin, H. Nakashima, L. Padgham, C. Sierra, and S. Willmott, editors, *Agentcities: Challenges in Open Agent Environments*, pages 25-31. Springer-Verlag, 2003.
8. Kim, W. (2001). *Knowledge-Based Personalization*, M.S. Thesis, Department of Computer Science, University of Georgia
9. Babaian, T. (2003). *Knowledge-Based Personalization*, in *Information Management: Support Systems and Multimedia Technology*, Idea Group Inc, Hershey, PA
10. Lieberman, H., Liu, H., Singh, P., Barry, B. (in press). *Beating some common sense into interactive applications*. AI Magazine.