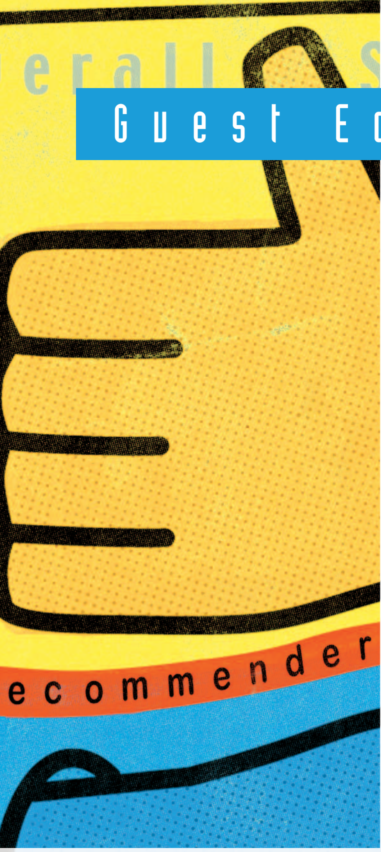

Recommender Systems

Alexander Felfernig and Gerhard Friedrich, *University of Klagenfurt*

Lars Schmidt-Thieme, *University of Hildesheim*

Recommender systems support users by identifying interesting products and services in situations where the number and complexity of offers outstrips the user's capability to survey them and reach a decision. Interest in recommender systems has dramatically increased owing to the demand for personalization technologies from large, successful e-commerce applications (such as Amazon.com). Nowadays, numerous online shops employ recommender applications, which many regard as a key enabling technology of e-commerce.

Corresponding applications recommend everything from news, Web sites, CDs, books, and movies to more complex items such as financial services, digital cameras, or e-government services. Recommendations are determined either by explicitly conducting sales dialogues with online users or by analyzing existing purchasing data from a single user or a community of users. Following the latter approach, the first recommender applications, developed in the mid-1990s, aimed to aggregate existing rating information to derive new user recommendations.

Recommendation approaches

One of the most frequently used representations of recommendation tasks is *collaborative filtering*,[1,2] essentially a nearest-neighbor method applied to a ratings matrix. For example, if two customers bought similar CDs and rated them similarly, the system would recommend to one customer CDs that the other customer bought and rated positively. Collaborative filtering perfectly implements the idea of word-of-mouth promotion, in which the opinions of friends and benchmarking reports are the predominant influences on a purchaser's buying decision.

The first applications of recommender technologies were the personalized recommendation of news and Web sites; the former application is often based on collaborative filtering, and the latter is based on *content-based filtering*.[3] Content-based filtering uses features of items the user liked in the past to infer new recommendations. Unlike collaborative approaches, content-based filtering can't provide serendipitous recommendations. It selects and recommends all products on the basis of purchasing information available from the current user. In both collaborative and content-based filtering, user profiles are long-term models. These approaches don't exploit deep knowledge about the product domain, and both are excellent for supporting the recommendation of simple products such as books or movies. A major strength of these approaches is that no additional knowledge acquisition efforts are necessary if historical data is available.

In recent years, researchers have proposed more complex statistical models for recommender systems that improve recommendation quality—for example, Bayesian networks with a hidden class variable[2] or compound classification models.[4] Some have tried hybridization strategies that combine collaborative and content-based information. Meanwhile, others have worked on attribute-aware recommender models that try to take co-user and product information into account in parallel.[5]

Customers purchasing complex products such as financial services, computers, or digital cameras need both information and intelligent interaction mechanisms that support the selection of appropriate solutions. So, an explicit representation of product, marketing, and sales knowledge is necessary. *Knowledge-based* approaches use this type of representation.[6–8] Such deep knowledge lets these recommender systems

1541-1672/07/$25.00 © 2007 IEEE
Published by the IEEE Computer Society

- calculate solutions that fulfill certain quality requirements,
- explain solutions to a customer, and
- support customers when the system can't find a solution.

In particular, explicit knowledge representation enables the validation of a recommender system's quality regarding calculated solutions. For example, in many applications, the generated solutions must suit customer requirements, adhere to legal regulations, and be in line with the company's sales strategy. Knowledge-based recommender technologies provide the formalisms necessary in this context. Unlike the word-of-mouth promotion that collaborative filtering provides, knowledge-based recommendation implements explicit sales dialogues that help users select items.

In this issue

The articles in this special issue cover a wide range of research topics. Francesco Ricci and Quang Nhat Nguyen present knowledge-based recommendation techniques that help on-the-move travelers select pretravel services. The recommendation environment relies on a critique-based approach that supports the acquisition and revision of user preferences. The recommendation system is connected to a recommender user interface specifically designed for mobile environments. In addition to the detailed discussion of the user interface and its support for revising critiques, the article also sketches the integration of a mobile recommendation environment (MobyRek) with a pretravel planning system (NutKing).

Jinhyung Cho, Kwiseok Kwon, and Yongtae Park propose a collaborative-filtering method based on group behavior theory in consumer psychology. The approach exploits the idea of defining dual information sources (groups of similar users and expert users). Depending on a user's receptivity regarding those groups, the system automatically adapts its recommendation strategy. This approach can improve classification accuracy compared to conventional collaborative filtering.

Product reviews and shared experiences about products are powerful information sources that can improve recommendation quality. Silvana Aciar, Debbie Zhang, Simeon Simoff, and John Debenham present a structured approach to including such information sources in calculating product recommendations. Their system exploits text-mining techniques to make textual information units applicable for calculating recommendations, which can improve the quality of recommendations based on collaborative filtering.

Gediminas Adomavicius and YoungOk Kwon present two new recommendation approaches that use multicriteria recommendation. Multicriteria ratings provide additional information about user preferences regarding several aspects of an item, such as story and acting quality in a movie. Experimental results from tests with a commercial data set show that multicriteria approaches improve recommendation accuracy compared to traditional single-rating strategies.

Publicly available recommender systems present a security problem. Attackers can introduce biased profiles to try to adapt the recommendation behavior in a way that's advantageous to the attacker. Bamshad Mobasher, Robin Burke, Runa Bhaumik, and J.J. Sandvig present an approach to understanding, identifying, and defeating such profile injection attacks. As a response to major attack patterns, they present algorithmic approaches for the design of more robust recommenders.

Neil J. Hurley, Michael P. O'Mahony, and Guenole C.M. Silvestre discuss security in recommender systems, focusing on a cost-benefit analysis of the financial gains that result from attacks on recommender systems. They also present a corresponding analysis framework.

Markus Zanker, Markus Jessenitschnig, Dietmar Jannach, and Sergiu Gordea evaluate different recommendation techniques using commercial data from an online cigar retailer. The evaluation compares three fundamentally different recommendation approaches (knowledge-based approaches, collaborative filtering, and content-based filtering). Primarily, they point out that collaborative-filtering algorithms provide good results for big data sets but are less effective for small data sets, where hybrid and knowledge-based approaches perform better.

Finally, Dan Melamed, Bracha Shapira, and Yuval Elovici present an approach to solving a problem that accompanies the application of collaborative-filtering systems: the free-ride problem. This occurs when users consume other people's evaluations without providing their own evaluations. The article presents an incentive mechanism based on a market-based model for pricing evaluations that motivates users to provide evaluations.

Future directions

In addition to the topics this special issue covers, we deem the following challenges as important future research directions.

Automated product data extraction

In a knowledge-based system, the quality of recommendations is directly correlated with the data quality. In many cases, product data are available only in an unstructured form. Furthermore, data in recommender knowledge bases become outdated. The major research focuses in this context are the automated extraction of product data from different information sources and the automated detection and adaptation of outdated product data. This includes identifying relevant information sources (for instance, in most cases, Web pages), extracting product data, and resolving contradictions in those data. A related recent challenge is extracting product information directly from digital multimedia products such as books, CDs, DVDs, and TV programs.

Community-based recommendation

State-of-the-art knowledge-based recommender environments focus on a single point of knowledge acquisition. Such centralized approaches can't guarantee the quality of knowledge bases because those knowledge

bases integrate the views of only a small group of experts. Mechanisms that effectively integrate the knowledge of a community of interest can potentially improve knowledge bases' quality and make related knowledge-acquisition processes more effective. Such participative architectures will enable the development of integrated knowledge bases and user applications.

Preference and assortment dynamics

With a few remarkable exceptions, most actual recommender models work on profiles that have been aggregated over time by simply collecting preference indicators. But in many domains, users' preferences develop over time, as do their assessments of certain products or product characteristics. For instance, a user who rated a digital camera as technically innovative two years ago might consider that camera outdated today. Taking this dynamic into account during model building will introduce a new level of complexity to recommender system models.

Intelligent testing

Successfully developing and maintaining recommender knowledge bases requires intelligent testing environments that can guarantee recommendations' correctness. So, future research must focus on developing mechanisms to automatically configure optimal test suites that both maximize the probability of identifying faulty elements in the recommender knowledge base and minimize the number of test cases needed to achieve this goal. Minimizing the number of test cases is important because domain experts must validate them manually.

Recommendation of configurable products and services

With the production of the Model T about 100 years ago, Henry Ford revolutionized manufacturing by employing mass production (the efficient production of many identical products). Nowadays, mass production is an outmoded business model, and companies must provide goods and services that fit customers' individual needs. In this context, *mass customization*—the production of highly variant products and services under mass production pricing conditions—has become the new paradigm.[9] A phenomenon accompanying mass customization is mass confusion,[10] which occurs when items are too numerous and complex for users to survey. Developing recommender technologies that apply to configurable products and services can help tackle mass confusion.

Context awareness

Integrating contextual elements into recommender algorithms is extremely important for improving recommendation results' quality. Such contextualization can play an important role in mobile environments—for instance, in the selection of travel services or destinations for on-the-move tourists. In such situations, recommendations depend not only on customer preferences but also on context, which can include attributes such as time of day, season, weather conditions, and ticket availability.

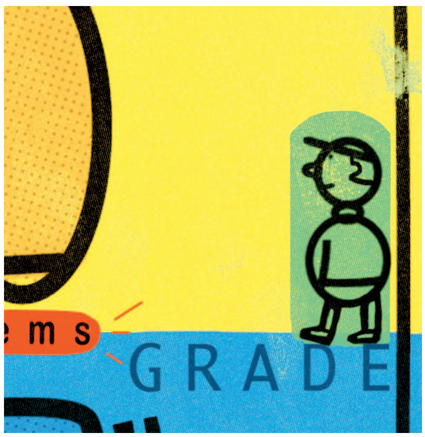

Theories of consumer buying behavior

Another major research issue is the active integration of theories from cognitive psychology and decision theory[11] into recommender applications. This integration will provide new insights into the major factors influencing consumer behavior and, consequently, could result in more efficient recommender systems in interactive selling environments.

Intelligibility and explanation

To be convincing, recommendations must be explained to customers. Only when they can challenge a recommendation and see why a system recommended a specific product will customers start to trust that system. Conceptual, knowledge-based, logical models that can provide such explanations should work in tandem with recommender systems based on machine-learning models. Combining statistical models' predictive power with the intelligibility of conceptual, human-understandable models might turn out to be a major success factor for recommender systems in the next decade.

Recommender technologies' high industrial demand has triggered the development of techniques to improve user interaction with recommenders (in terms of more user-friendly and effective dialogue structures), the quality of result sets presented to users (in terms of precision/recall), and user trust in presented recommendations. The articles presented in this special issue address these techniques, and the field presents a wide array of open research questions.

Acknowledgments

We thank all the authors for their contributions to this special issue. Furthermore, we thank the reviewers for their valuable comments, which significantly contributed to the accepted articles' high quality. Finally, we express our gratitude to Editor in Chief James Hendler and the *IEEE Intelligent Systems* staff, who supported this special issue on recommender systems and served as mentors and guides for this project.

References

1. J. Konstan et al., "GroupLens: Applying Collaborative Filtering to Usenet News," *Comm. ACM*, vol. 40, no. 3, 1997, pp. 77–87.
2. T. Hofmann, "Latent Semantic Models for Collaborative Filtering," *ACM Trans. Information Systems*, vol. 22, no. 1, 2004, pp. 89–115.
3. M. Pazzani and D. Billsus, "Learning and Revising User Profiles: The Identification of Interesting Web Sites," *Machine Learning*, vol. 27, no. 3, 1997, pp. 313–331.
4. L. Schmidt-Thieme, "Compound Classification Models for Recommender Systems," *Proc. IEEE Int'l Conf. Data Mining* (ICDM 05), IEEE Press, 2005, pp. 387–385.
5. K. Tso and L. Schmidt-Thieme, "Evaluation of Attribute-Aware Recommender System Algorithms on Data with Varying Characteristics," *Proc. 10th Pacific-Asia Conf. Knowledge Discovery and Data Mining* (PAKDD 06), Springer, 2006, pp. 831–840.
6. R. Burke, "Knowledge-Based Recommender Systems," *Encyclopedia of Library and Infor-*

The Authors

Alexander Felfernig is an associate professor at the University Klagenfurt's Institute of Applied Informatics. He's also a member of the university's Intelligent Systems and Business Informatics research group and cofounder and CEO of ConfigWorks. He received his PhD in computer science from the University of Klagenfurt. His research interests include knowledge-based configuration, knowledge-based recommenders in e-commerce, psychological aspects of consumer buying behavior, model-based diagnosis with a special focus on knowledge acquisition, and approaches to test automation in knowledge-based-systems development. He's a member of the IEEE Computer Society. Contact him at Alpen-Adria Universität Klagenfurt, Universitätsstrasse 65, 9020 Klagenfurt, Austria; alexander.felfernig@uni-klu.ac.at.

Gerhard Friedrich is a professor at the University of Klagenfurt, where he directs the research group on Intelligent Systems and Business Informatics. His research interests include knowledge-based systems for configuring products and services, knowledge-based advisory systems, distributed knowledge-based systems, knowledge acquisition, integration of knowledge-based systems, diagnosis and repair of processes, and self-healing Web services. He received his PhD in computer science from the Technical University Vienna. He was previously head of the technology center for configurators and diagnosis systems of Siemens Austria, where under his responsibility numerous knowledge-based systems were successfully implemented. He's an editor of *AI Communications* and is associate editor of the *International Journal of Mass Customisation*. Contact him at Alpen-Adria Universität Klagenfurt, Universitätsstrasse 65, 9020 Klagenfurt, Austria; gerhard.friedrich@uni-klu.ac.at.

Lars Schmidt-Thieme is a professor at the Institute for Economics and Information Systems and the Institute of Computer Science at the University of Hildesheim, where he directs the Information Systems and Machine Learning Lab. His research interests include machine learning and data mining, especially classification and pattern mining for complex data. He received his PhD in economics from the University of Karlsruhe. He's involved in several research projects, such as the EC FP6 integrated project X-Media. Contact him at the Univ. of Hildesheim, Marienburger Platz 22, D-31141 Hildesheim, Germany; schmidt-thieme@ismll.uni-hildesheim.de.

mation Systems, vol. 69, supplement 32, A. Kent, ed., Marcel Dekker, 2000, pp. 180–200.

7. A. Felfernig et al., "An Environment for the Development of Knowledge-Based Recommender Applications," *Int'l J. Electronic Commerce*, vol. 11, no. 2, 2006, pp. 11–34.
8. A. Felfernig and B. Gula, "An Empirical Study on Consumer Behavior in the Interaction with Knowledge-Based Recommender Applications," *Proc. IEEE Joint Conf. E-Commerce Technology and Enterprise Computing, E-Commerce and E-Services* (CEC 06/EEE 06), IEEE Press, 2006, pp. 288–296.
9. B. Pine and S. Davis, *Mass Customization: The New Frontier in Business Competition*, Harvard Business School Press, 1999.
10. C. Huffman and B. Kahn, "Variety for Sale: Mass Customization or Mass Confusion," *J. Retailing*, vol. 74, no. 4, 1998, pp. 491–513.
11. J. Payne, J. Bettman, and E. Johnson, *The Adaptive Decision Maker*, Cambridge Univ. Press, 1993.