Introduction to the IEEE Intelligent Systems Special Issue: Recommender Systems

Alexander Felfernig
Gerhard Friedrich
Lars Schmidt-Thieme

Recommender systems support users in the identification of interesting products and services in situations where the amount and complexity of offers outstrips the capability of a user to survey it and to reach a decision. The interest in recommender systems has dramatically increased owing to the demand for personalization technologies by large, successful e-commerce applications (for example, see www.amazon.com). Nowadays, recommender applications are deployed on many online shops and are regarded as one of the key enabling technologies of e-commerce. The corresponding applications range from the recommendation of news, web sites, compact disks, books, or movies to the recommendation of more complex items such as financial services, digital cameras, or e-government services. The recommendations are determined either by explicitly conducting sales dialogs with online users or by analyzing existing purchasing data of a single user or a community of users. Following the latter approach, the focus of the first recommender applications developed in the mid-1990s was to aggregate existing rating information in order to derive new recommendations for the current user.

One of the most frequently used representation of recommendation tasks is *collaborative filtering* which is in essence a nearest-neighbour method applied to a matrix of ratings: if two customers have bought similar compact disks (CDs) in the past and have rated those CDs in a similar way, CDs (with a positive rating) bought by only one of them, are recommended to the other one. Collaborative filtering perfectly implements the idea of word-of-mouth promotion where a buying decision is predominantly influenced by the opinions of friends and benchmarking reports. First applications of recommender technologies were the personalized recommendation of news and web sites where the former application is often based on collaborative filtering and the latter is based on a *content-based filtering* approach. Content-based filtering is a special type of information filtering that uses features of items the user has liked in the past to infer new recommendations. In contrast to collaborative approaches, content-based filtering is limited in that it cannot provide serendipitous recommendations. It selects and recommends all products based on purchasing information available from the current user. In the case of collaborative filtering as well as in the case of content-based filtering, user profiles are long-term models. Both approaches do not exploit deep knowledge about the product domain and are excellent techniques supporting recommendation processes for simple products such as books or movies. A major strength of these approaches is that no additional knowledge acquisition efforts are needed provided that historical data is available.

Besides simple collaborative and content-based filtering, in recent years many more complex statistical models for recommender systems have been proposed that improve recommendation quality, e.g., Bayesian networks with a hidden class variable or compound classification models. Collaborative and content-based information early has been tried to be combined with so-called hybridization strategies. Meanwhile there is also research on attribute-aware recommender models that try to take information about co-users and about products into account in parallel.

In contrast to model-based recommender systems, knowledge-based approaches make use of an explicit representation of product, marketing and sales knowledge. Compared to customers purchasing simple products such as books or compact disks, customers purchasing more complex products such as financial services, computers, or digital cameras are much more in the need of information and in the need of intelligent interaction mechanisms which support the selection of appropriate solutions. Therefore, an explicit representation of product, marketing and sales knowledge is needed. Such deep knowledge allows (a) the calculation of solutions which fulfil certain quality requirements, (b) the explanation of solutions to a customer, and (c) the support of customers in situations where no solution can be found. In particular, explicit knowledge representation enables the validation of the quality of a recommender system regarding calculated solutions. E.g., in many applications the generated solutions have to be in line with the company's sales strategy, suite the requirements of customers, and adhere to legal regulations. Primarily, knowledge-based recommender technologies provide the formalisms that are needed in this context. In contrast to word-of-mouth promotion implemented by collaborative filtering, knowledge-based recommendation implements explicit sales dialogs which support users in the item selection process.

The papers contained in this special issue on recommender systems comprise a wide range of research topics. *Ricci* and *Nguyen* present knowledge-based recommendation techniques which support on-the-move travellers in the selection of desired travel services. The recommendation environment relies on a critiquing-based approach which supports the acquisition and revision of user preferences. The recommendation engine is connected with a recommender user interface which is specifically designed for the application in mobile environments. In addition to the detailed discussion of the user interface and the supported tweaking critiquing techniques, the article as well sketches the integration of the mobile recommendation environment (MobyRek) with the pre-travel planning system (NutKing).

Cho, *Kwon*, and *Park* propose a novel collaborative filtering method which is based on group behaviour theory in consumer psychology. The approach exploits the idea of defining dual information sources (groups of similar users and groups of expert users). Depending on the degree of a user's receptivity regarding those groups, the recommender system automatically adapts its recommendation strategy. This new collaborative filtering (CF) approach can improve classification accuracy compared to conventional CF.

Product reviews and shared experiences about products are powerful information sources which can be utilized to improve the quality of recommendations. *Aciar*, *Zhang*, *Simoff*, and *Debenham* present a structured approach to include such information sources into the calculation of product recommendations. In this context, text mining techniques are exploited in order to make textual information units applicable for calculating recommendations. Thus, the quality of recommendations based on collaborative filtering can be improved.

Adomavicius and *Kwon* present two new recommendation approaches which use multi-criteria recommendation. Multi-criteria ratings provide additional information about the preferences of a user regarding several aspects of an item, e.g., quality of the story and acting in a movie. Experimental results from tests with a commercial dataset show that multi-criteria recommendation approaches improve recommendation accuracy if compared to traditional single rating strategies.

Publicly available recommender systems present a security problem. Attackers may introduce biased profiles with the goal to adapt the recommendation behaviour in a way which is advantageous to the attacker. *Mobasher*, *Burke*, *Bhaumik*, and *Sandvig* present an approach to understand, identify, and defeat such profile injection attacks. The response to major attack patterns are algorithmic approaches for the design of more robust recommenders.

The issue of security in recommender systems is additionally discussed by *Mahony*, *Hurley*, and *Silvestre* who focus in their paper on a cost-benefit analysis of the financial gains that are realised by attacks on recommender systems. A corresponding analysis framework is presented in the paper.

The paper of *Zanker*, *Jessenitschnig*, *Jannach*, and *Gordea* presents the evaluation of different recommendation techniques on the basis of a commercial dataset from the area of fine cigars. The evaluation compares three fundamentally different recommendation approaches (knowledge-based recommendation vs. collaborative and content-based filtering). Primarily, it is pointed out that CF algorithms provide good results for big datasets but are less effective for small datasets, where hybrid and knowledge-based approaches perform better.

Finally, the article of *Melamed*, *Shapira*, and *Elovici* presents an approach to solve the free-ride problem (users consume item evaluations provided by others without providing their own evaluations) coming along with the application of collaborative filtering systems. A new incentive mechanism motivating users to provide evaluations is presented which is based on a market-based model for pricing evaluations.

Beside the research topics discussed in the papers of this special issue, we deem the following challenges as important future research directions.

Automated product data extraction. Knowledge-based recommenders exploit deep product, marketing, and sales knowledge for determining recommendations. The quality of a recommendation is directly correlated with the quality of the corresponding product data. In many cases, product data are only available in an unstructured form, furthermore, existing product data in recommender knowledge bases become outdated. The major research focus in this context is the automated extraction of product data from different information sources and the automated detection and adaptation of outdated product data. This includes the identification of relevant information sources (e.g., in most cases Web pages), the extraction of product data and the resolution of contradictions in extracted product data. Also extracting product information directly from digital multimedia products such as books, CDs, DVDs, TV programs etc. is a challenging recent problem.

Community-based Recommendation. State-of-the-art knowledge-based recommender environments focus on a single point of knowledge acquisition. Such centralized approaches can not guarantee the optimal quality of knowledge bases since only the view of a small group of experts is integrated into the recommendation knowledge base. Mechanisms which allow an effective integration of the knowledge of a community of interest promise high potential for improving the quality of knowledge bases and the effectiveness of related knowledge acquisition processes. Such participative architectures will enable the development of integrated knowledge bases and applications to users.

Preference and assortment dynamics: with a few remarkable exceptions, most actual recommender models work on profiles that have been aggregated over time by simple extension. But in many domains, preferences of users develop and the assessment of certain products or product characteristics changes over time – a digital camera that has been technically innovative two years ago, will be outdated today. Taking this dynamics systematically into account during model building will introduce a new level of complexity into recommender system models in the next years.

Intelligent Testing. The major precondition for successfully developing and maintaining recommender knowledge bases are intelligent testing environments which can guarantee the correctness of recommendations. In this context, mechanisms have to be developed which automatically configure optimal test suites which maximize the probability of identifying faulty elements in the recommender knowledge base and at the same time minimize the number of test cases needed to achieve this goal. The minimization of test cases is an important issue in this context, since the validation of test cases is a manual process to be conducted by domain experts.

Recommendation of Configurable Products and Services. With the production of the T-model about hundred years ago, Henry Ford revolutionized manufacturing processes by introducing the concepts of Mass Production (efficient production of a high number of identical products). Nowadays, Mass Production is a business model of the past and companies are forced to provide goods and services which fit the individual needs of a customer. In this context, Mass Customization has been established as a new paradigm which is defined as the production of highly variant products and services under Mass Production pricing conditions. A phenomenon which comes along with the establishment of Mass Customization concepts is Mass Confusion. This phenomenon occurs in situations where the amount and complexity of items offered outstrips the capability of a user to survey it. Recommender systems can be key-enabling technologies tackling the challenges of Mass Confusion. A precondition for reaching this goal is the development of recommender technologies as well applicable for configurable products and services.

Context Awareness. Integrating contextual elements into recommender algorithms is extremely important for improving the quality of recommendation results. Such a 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 do not only depend on the personal preferences of customers but as well on the context which can be described by attributes such as time of the day, season, weather conditions, availability of tickets, etc.

Theories of Consumer Buying Behaviour. A major research issue is the active integration of different psychological theories from the area of cognitive psychology and decision theory into recommender applications. This integration will provide new insights on the major influencing factors of consumer buying behaviour and consequently more efficient recommender systems in interactive selling environments.

Intelligibility and Explanations. To be convincing, recommendations must be explained to customers. Only when they have a chance to challenge a specific recommendation and see, why a specific product has been recommended to them, customers will start to trust a system. Recommender systems based on Machine Learning models will have to be complemented by more conceptual, knowledge-based, logical models that will be able to provide such explanations. To combine the predictive power of the statistical models with the intelligibility of conceptual, human understandable models may turn out to be one of the major success factors for recommender systems in the next decade.

We want to thank all the authors for their contributions to this special issue on recommender systems. Furthermore, we want to thank the reviewers for their valuable comments which significantly contributed to the high quality of the accepted papers. Finally, we express our gratitude to Jim Hendler, Dennis Taylor, and Hilda Carman who supported this special issue on recommender systems and served as mentor and guide for this project.

Alexander Felfernig
Gerhard Friedrich
Lars Schmidt-Thieme

References

[1] R. Burke. Knowledge-based Recommender Systems. *Encyclopedia of Library and Information Systems*, 69(32), 2000.

[2] A. Felfernig, G. Friedrich, D. Jannach, and M. Zanker, An Environment for the Development of Knowledge-based Recommender Applications, *International Journal of Electronic Commerce* (*IJEC*), 11(2): 11-34, 2006.

[3] A. Felfernig and B. Gula, An Empirical Study on Consumer Behavior in the Interaction with Knowledge-based Recommender Applications, in IEEE Joint Conference on E-Commerce Technology (CEC06) and Enterprise Computing, E-Commerce and E-Services (EEE06). San Francisco, California: IEEE, 2006, pp. 288–296 (best papers).

[4] C. Huffman and B. Kahn, Variety for Sale: Mass Customization or Mass Confusion, *Journal of Retailing*, (74): 491–513, 1998.

[5] Th. Hofmann, Latent Semantic Models for Collaborative Filtering, *JACM Transactions on Information Systems 22/1: 89-115, 2004*.

[6] J. Konstan, B. Miller, D. Maltz, J. Herlocker, L. Gordon, and J. Riedl. GroupLens: applying collaborative filtering to Usenet news Full text. *Communications of the ACM*, 40(3): 77-87, 1997.

[7] J. Payne, J. Bettman, and E. Johnson. The Adaptive Decision Maker. *Cambridge University Press*, 1993.

[8] M. Pazzani and D. Billsus. Learning and Revising User Profiles: The Identification of Interesting Web Sites. *Machine Learning*, (27): 313–331, 1997.

[9] B. Pine and S. Davis, Mass Customization: The New Frontier in Business Competition. *Harvard Business School Press*, 1999.

[10] K. Tso and L. Schmidt-Thieme, Evaluation of Attribute-aware Recommender System Algorithms on Data with Varying Characteristics, *Proceedings of the Tenth Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, Singapore, 2006, pp. 831-840.

[11] L. Schmidt-Thieme, Compound Classification Models for Recommender Systems, *Proceedings of the IEEE International Conference on Data Mining* (*ICDM'05*), 2005, Houston, USA, pp. 387-385.