Preface to the special issue on context-aware recommender systems

Gediminas Adomavicius · Dietmar Jannach

1 Context-aware recommender systems

Recommender systems represent a popular area of personalization technologies that has enjoyed a tremendous amount of research and development activity in both academia and industry in the last 10–15 years. Recommender systems research typically explores and develops techniques and applications for recommending various products or services to individual users based on the knowledge of users' tastes and preferences as well as users' past activities (such as previous purchases), which are applicable in a variety of domains and settings (Jannach et al. 2010).

While a substantial amount of research has already been performed in the area of recommender systems, many existing approaches have focused on recommending the most relevant items to users and do not take into account any additional contextual information, such as time, location, weather, the user's current goals, the user's mood, presence of other people, or the type of the device through which the recommendation is consumed. In other words, traditionally recommender systems deal with applications having only two types of entities, users and items, and do not put them into a context when providing recommendations. However, the importance of contextual information has been recognized by researchers and practitioners in many

G. Adomavicius
Department of Information and Decision Sciences, Carlson School of Management,
University of Minnesota, Minneapolis, MN, USA
e-mail: gedas@umn.edu

D. Jannach (✉)
Department of Computer Science, TU Dortmund, Dortmund, Germany
e-mail: dietmar.jannach@tu-dortmund.de

disciplines (such as ubiquitous and mobile computing, e-commerce, marketing, information retrieval, and many others), and, in the past several years, context-awareness has been increasingly recognized as a critical issue in many recommendation applications and has been explored by a number of recommender systems researchers. This is evidenced by an increasing number of papers on context-aware recommender systems (CARS) that appear in major conferences (such as the *ACM Conference on Recommender Systems*) as well as the continued success of several events specifically dedicated to context-aware recommender systems (such as the annual *Workshop on Context-Aware Recommender Systems* in 2009, 2010, 2011 and 2012, *the Workshop on Context-Awareness in Retrieval and Recommendation* in 2011 and 2012, and the *Challenge on Context-Aware Movie Recommendation* in 2010 and 2011).

A recent survey on context-aware recommender systems as well as the discussion of the potential directions for future research can be found in Adomavicius and Tuzhilin (2011). Looking at the current state of the art in context-aware recommender systems, we can broadly classify the main research issues, challenges, and directions into the following four general categories:

1. **Fundamentals**, i.e., understanding the notion of context and modeling context in recommender systems.
2. **Algorithms**, i.e., developing recommendation algorithms that can incorporate contextual information into recommender systems in advantageous ways.
3. **Evaluation**, i.e., in-depth evaluation of context-aware recommender systems' performance, their benefits and limitations.
4. **Engineering**, i.e., designing general-purpose architectures, frameworks, and approaches to facilitate the development, implementation, deployment, and use of context-aware recommendation capabilities.

A lot of existing research on context-aware recommender systems can be classified under the "Algorithms" category. In particular, several recommendation techniques have been proposed that take advantage of contextual information using one of the following algorithmic paradigms (Adomavicius and Tuzhilin 2011): (a) *contextual pre-filtering*, where contextual information is used to select (or construct) a subset of data relevant for that specific context (i.e., input data is contextualized), after which any traditional (non-context-aware) recommendation algorithm can be applied on the selected data to compute recommendations; (b) *contextual post-filtering*, where contextual information is initially ignored and any traditional (non-contextual) recommendation algorithm can be applied on the entire dataset, after which the resulting set of recommendations is adjusted (contextualized) for each user using the contextual information; and (c) *contextual modeling*, where contextual information is used directly in the modeling technique as part of the data-driven learning process. In addition, some other recommender systems can also make recommendations by using *context-driven querying and search*, where contextual information (such as the current user's location or current time) can be simply used as a query to retrieve relevant information as recommendations (e.g., a list of nearby museums that are currently open). As compared to the "Algorithms" category, the other three aforementioned categories have been relatively underexplored.

This UMUAI special issue on context-aware recommender systems builds on the recent activities in this area and provides important and interesting contributions across a wide number of the aforementioned research categories. In particular, this special issue includes four papers, where the first paper falls under "Algorithms", the second and third papers under "Evaluation", and the fourth paper under "Engineering". Brief summaries of the papers are provided in the next section. We believe that context-aware recommender systems represent an increasingly active and highly problem-rich research area, and hope that the papers in this special issue will spark many further developments not only in "Algorithms", "Evaluation", and "Engineering", but in "Fundamentals" as well. Clearly, *context* is a complex notion, and there are many different approaches of conceptualizing context in recommender systems. Some simple questions that are yet to be fully explored include Adomavicius et al. (2011): Should the contextual dimensions of interest be modeled as static, or can they change over time? Or, should contextual information be assumed to be fully observable to the recommender system or can it be latent? The majority of existing research on context-aware recommender systems follows the assumption that contextual dimensions are stable and observable, but implications of other modeling assumptions and approaches should also be explored.

2 Papers in this issue

The paper "Experimental Evaluation of Context-Dependent Collaborative Filtering Using Item Splitting" (Baltrunas and Ricci 2014) by Linas Baltrunas and Francesco Ricci is concerned with algorithmic approaches for the development of CARS and provides an in-depth analysis of the "item splitting" pre-filtering method for context-aware recommendation. As usual in CARS, the main idea of their technique is that the same recommendable item can be experienced or consumed by users in different contextual settings (e.g., in summer or in winter) and that the user preferences (in the form of user-provided ratings) for the items may correspondingly depend on this contextual setting. In the case when there is sufficient statistical evidence that users actually rate some item differently in certain contexts, the authors propose to split the item into two or theoretically even more fictitious items, one for each contextual condition. Afterwards, in the prediction phase, the proposed recommendation technique then takes the user's current contextual condition into account and returns the rating prediction for the fictitious item that matches the users's current context. An evaluation of their method on smaller real-world and larger semi-synthetic data sets shows that their context-aware item splitting approach can help to improve the prediction quality of a recommender system in the analyzed application scenarios.

The paper "Comparing Context-Aware Recommender Systems in Terms of Accuracy and Diversity: Which Contextual Modeling, Pre-Filtering and Post-Filtering Methods Perform the Best" (Panniello et al. 2014) by Umberto Panniello, Alexander Tuzhilin, and Michele Gorgoglione is devoted to a systematic evaluation and comparison of the effectiveness of existing techniques for building context-aware recommendations with respect to different evaluation dimensions. Their analysis takes into account different recommendation tasks ("find-all" vs. "top-k"), two utility metrics

(accuracy vs. diversity), the granularity of the processed contextual information, as well as other characteristics of the underlying real world data sets. Among the findings, their analysis shows that the accuracy of all tested CARS techniques decreases when the context information has a finer granularity and less information is available for each contextual situation. This can be seen as an indicator that still better techniques are required to more effectively recommend items that match the precise contextual situation of the user. Overall, the comparison revealed that there is no clear winning technique that dominates the others with respect to all evaluation dimensions. The Contextual Modeling (CM) and Post-Filtering approaches in many situations perform well in terms of accuracy; on the other hand, Exact Pre-Filtering often leads to more diverse recommendations. The usual trade-off between accuracy and diversity however also exists in context-aware recommenders. The results of the study finally indicate that CM approaches are comparatively good at balancing this trade-off in general.

The paper "Time-Aware Recommender Systems: A Comprehensive Survey and Analysis of Existing Evaluation Protocols" (Campos et al. 2014) by Pedro Campos, Fernando Díez, and Iván Cantador focuses on "time" as one of the most valuable contextual factors in many RS domains. They survey and classify various existing approaches to modeling time aspects in the user-modeling and recommendation phases. In particular, the authors review common evaluation practices and shed light on methodological issues related to the comparative evaluation of time-aware recommender systems based on historical data sets. They demonstrate that the choice of the evaluation conditions impacts the ranking of different recommendation strategies, which to some extent can be observed in context-unaware recommender systems, too. To overcome such issues, the authors propose a set of guidelines and a corresponding methodological framework for a robust and fair evaluation process. The work, therefore, represents an important step in the direction of improved, standardized and, thus, more reproducible evaluation procedures for an important class of CARS.

Finally, the paper "Hybreed: A Software Framework for Developing Context-Aware Hybrid Recommender Systems" (Hussein et al. 2014) by Tim Hussein, Timm Linder, Werner Gaulke, and Jürgen Ziegler addresses the systems engineering perspective and introduces a software framework for the development of context-aware and hybrid recommenders. Despite the high practical relevance of recommender systems in industry, little research on engineering aspects of such systems has been done, and no comprehensive software framework is yet available that supports component-based development approaches for such complex systems. One of the challenges in the design of such a framework lies in the choice of the appropriate level of abstraction so that the individual components can be (re)used in different application domains, and a number of possible data sources can be easily integrated. In particular, the question of how to acquire and process various types of contextual information (as done in the more general field of context-aware computing) is identified as an important requirement for the design of such a framework. The paper introduces the Hybreed framework and discusses its underlying Dynamic Contextualization approach, the general architecture, as well as the functionality of the pre-implemented framework components and presents the first analytical and empirical evaluation of the framework.

References

Adomavicius, G., Mobasher, B., Ricci, F., Tuzhilin, A.: Context-aware recommender systems. AI Mag. **32**(3), 67–80 (2011)

Adomavicius, G., Tuzhilin, A.: Context-aware recommender systems. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P.B. (eds.) Recommender Systems Handbook, pp. 217–253. Springer, New York (2011)

Baltrunas, L., Ricci, F.: Experimental evaluation of context-dependent collaborative filtering using item splitting. User Model User-Adap. Inter. (2014). doi:10.1007/S11257-012-9137-9

Campos, P.G., Díez, F., Cantador, I.: Time-aware recommender systems: A comprehensive survey and analysis of existing evaluation protocols. User Modeling and User Model User-Adap. Inter. (2014). doi:10.1007/S11257-012-9136-x

Hussein, T., Linder, T., Gaulke, W., Ziegler, J.: Hybreed: A software framework for developing context-aware hybrid recommender systems. User Modeling and User Model User-Adap. Inter. (2014). doi:10.1007/S11257-012-9134-z

Jannach, D., Zanker, M., Felfernig, A., Friedrich, G.: Recommender Systems: An Introduction. Cambridge University Press, Cambridge (2010)

Panniello, U., Tuzhilin, A., Gorgoglione, M.: Comparing context-aware recommender systems in terms of accuracy and diversity: Which contextual modeling, pre-filtering and post-filtering methods perform the best. User Model User-Adap. Inter. (2014). doi:10.1007/S11257-012-9135-y

Author Biographies

Gediminas Adomavicius is a professor of information and decision sciences at the University of Minnesota. He received his Ph.D. in Computer Science from New York University. His research focuses on personalization technologies, recommender systems, data mining, and complex electronic market mechanisms. His research has been published in a number of leading computer science and information systems journals, including *IEEE Transactions on Knowledge and Data Engineering*, *ACM Transactions on Information Systems*, *Data Mining and Knowledge Discovery*, *IEEE Intelligent Systems*, *Information Systems Research*, *Management Science*, *MIS Quarterly*, and *INFORMS Journal on Computing*. He serves on editorial boards of *Information Systems Research* and *INFORMS Journal on Computing*.

Dietmar Jannach is a professor of computer science at TU Dortmund, Germany, and chair of the department's e-Services Research Group. His main research interests lie in the application of artificial intelligence and knowledge-based systems technology to real-world problems in particular in e-business environments. His research works have been published in journals such as *Artificial Intelligence*, *IEEE Transactions on Knowledge and Data Engineering*, *Knowledge and Information Systems*, the *Journal of Web Semantics*, *IEEE Intelligent Systems* or *AI Magazine* and at conferences such as *IJCAI*, *ICDM* or *ECAI*. Dietmar Jannach was also one of the co-founders of ConfigWorks GmbH, a company focusing on next-generation interactive recommendation and sales advisory systems.