User Model User-Adap Inter
DOI 10.1007/s11257-007-9042-9

ORIGINAL PAPER

Mediation of user models for enhanced personalization in recommender systems

Shlomo Berkovsky · Tsvi Kuflik · Francesco Ricci

Received: 21 March 2007 / Revised: 21 July 2007 / Accepted in revised form: 12 September 2007

Abstract Provision of personalized recommendations to users requires accurate modeling of their interests and needs. This work proposes a general framework and specific methodologies for enhancing the accuracy of user modeling in recommender systems by importing and integrating data collected by other recommender systems. Such a process is defined as user models mediation. The work discusses the details of such a generic user modeling mediation framework. It provides a generic user modeling data representation model, demonstrates its compatibility with existing recommendation techniques, and discusses the general steps of the mediation. Specifically, four major types of mediation are presented: cross-user, cross-item, cross-context, and cross-representation. Finally, the work reports the application of the mediation framework and illustrates it with practical mediation scenarios. Evaluations of these scenarios demonstrate the potential benefits of user modeling data mediation, as in certain conditions it allows improving the quality of the recommendations provided to the users.

Keywords Recommender systems · Ubiquitous user modeling · Mediation of user modeling data

S. Berkovsky (✉)
Department of Computer Science and Software Engineering, University of Melbourne, Melbourne, VIC 3010, Australia
e-mail: shlomo@csse.unimelb.edu.au

S. Berkovsky · T. Kuflik
University of Haifa, Haifa, Israel

T. Kuflik
e-mail: tsvikak@is.haifa.ac.il

F. Ricci
Free University of Bozen-Bolzano, Bolzano, Italy
e-mail: fricci@unibz.it

1 Introduction

During the last decade, the quantity of potentially interesting products or information services available online has been growing rapidly and now exceeds human processing capabilities (Maes 1994). Moreover, there are many information search situations where the users would like to choose among a set of alternative items or services, but do not have sufficient knowledge, capabilities or time to make such decisions. As such, there is a pressing need for intelligent systems that advise users while taking into account their personal needs and interests. Such systems can deliver tailored service in a way that will be most appropriate and valuable to the users. This type of systems is referred to in the literature as personalization systems (Mulvenna et al. 2000).

This work focuses on recommender systems (Resnick and Varian 1997), a typical example of personalization systems. Recommender systems provide users with recommendations about products and services they may like. They generate personalized recommendations, i.e., recommendations that are tailored to the user. This task is achieved by exploiting various knowledge sources, which store information collected during past interactions with users searching or providing recommendations, and the evaluations of those recommendations. Extensive research in recommender systems started over a decade ago and yielded a wide variety of recommendation techniques, such as content-based filtering (Morita and Shinoda 1994), collaborative filtering (Resnick et al. 1994), knowledge-based recommendation (Burke 2000), utility-based recommendation (Manouselis and Sampson 2004) and their multiple hybridizations (Burke 2002). These techniques are widely discussed in the literature and in several surveys of the state-of-the-art recommender systems (Adomavicius and Tuzhilin 2005a; Montaner et al. 2003; Schafer et al. 2000).

Whatever the specific technology exploited by a recommender system, it can provide high quality recommendations to users only after having modeled their preferences. This information is typically referred to in the literature as the User Model (UM) (Kobsa 2001). The task of collecting user modeling data is typically performed in two ways: (1) explicit—through provision of the required information explicitly by the user, or (2) implicit—through applying various reasoning mechanisms that infer the required information based on the user's observable behavior (Hanani et al. 2001). The explicit collection of user modeling data is considered to be an accurate, but time- and effort-consuming task, typically avoided by the user. Alternatively, the implicit collection involves automated reasoning mechanisms, which can misinterpret user behavior. In practice, the explicit and implicit approaches may also be combined (Pazzani 1999).

In general, the quality of the recommendations provided to the user depends largely on the characteristics of the UM, e.g., how accurate it is, what amount of information it stores, and whether this information is up to date. Hence, as a general rule, the more information is stored in the UM, i.e., the more knowledge the system has obtained about the user, the better the quality of the recommendations will be. In this context, quality refers to the capability of the system to suggest exactly those products or services that the user will select and purchase, or to predict correctly those items that the user would like. In practice, obtaining sufficient user modeling data to deliver high quality recommendations is difficult. This is especially important at the initial stages of the interaction with the user, when little information about the user is available. At these stages, all the existing recommendation techniques face the bootstrapping problem, i.e., a situation where the available information about the user and/or items does not suffice to provide high quality recommendations (Linden et al. 2003).

When analyzing current recommender systems, one can see that, typically, every system builds and maintains a proprietary collection of UMs (Montaner et al. 2003). Practically, this means that the collected user modeling data are tailored to: (1) the specific content (products or products categories) offered by the recommender system, e.g., movies (Good et al. 1999), music (Aguzzoli et al. 2002), news items (Claypool et al. 1999), tourism (Ricci et al. 2002), and so on, and (2) the recommendation technique being exploited by the system, e.g., collaborative filtering (Herlocker et al. 1999), content-based filtering (Morita and Shinoda 1994), demographic filtering (Krulwich 1997), or some of their hybridizations (Burke 2002). Thus, a large amount of heterogeneous (and possibly overlapping) user modeling data are scattered among various systems. However, practical recommender systems (and, especially commercial ones) neither allow other external recommender systems to access them, nor share their proprietary user modeling data. Despite this, it is reasonable to hypothesize that recommender systems could potentially benefit from enriching their user modeling data by importing and integrating user modeling data collected by other recommender systems, and therefore provide better recommendations to the users.

User modeling data integration can be achieved through a process that is referred to in this work as the *mediation* of UMs (Berkovsky 2006). Mediation of UMs is a process of importing and integrating the user modeling data collected by other recommender systems for the purposes of a specific recommendation task. Hence, the primary goal of the mediation is to instantiate UMs through inferring the required user modeling data from other data imported from other systems. The mediation enriches the existing UMs (or bootstraps empty UMs) in the target recommender systems and, as a result, facilitates provision of better recommendations.

Let us introduce an example mediation scenario in the realm of digital entertainment. Consider a Web-based network of sites providing personalized entertainment recommendations. The network includes music, movies, TV programs, books, and humor recommender systems. Also, consider a user requesting a personalized movie recommendation from the movies recommender system. Although the movies recommender system collected certain user modeling data about the user in the past, these may not be sufficient for providing high quality recommendations. To enhance the recommendations, the movies recommender system can obtain a more accurate UM by importing and integrating UMs collected by other systems. For example, the user's favorite art genres can be imported from the UMs generated for recommendation of TV programs and/or books recommender systems, while the user's favorite composers can be imported from the music recommender system. These data can be used, for example, for recommending a movie of the favorite genre, where the music was composed by the favorite composer.

The idea of UM mediation brings forward a number of challenges. The first one refers to the nature of the information market, and its business models. Due to commercial competition, practical real-life recommender systems currently do not cooperate, and do not share their user modeling data. In Adomavicius and Tuzhilin (2005b), the authors point out that typical recommender systems are either provider-centric (i.e., each provider has its own recommendation engine to tailor its content to consumers) or market-centric (i.e., providing recommendations for a specific marketplace in a particular industry or sector). The authors claim that the lack of technical data sharing solutions in the existing recommender systems is mainly explained by business limitations imposed on the exchange of user-related information among competing parties in the same market.

The second challenge refers to guaranteeing users' privacy. UMs collected by a certain recommender system may contain private and sensitive information that the users would not like to be disclosed to other systems, and possibly to untrusted parties (Cranor et al. 1999). For this reason, many recommender services that store sensitive information about their users declare in their privacy policies that no personal information stored by the system will be transferred to other parties (Wang and Kobsa 2006). As a result, they are committed not to transfer information stored in their UMs to other systems.

The third challenge refers to practical and technical considerations of the mediation. For example, in a distributed setting various recommender systems need to connect to each other through slow, inherently unreliable, and error-prone communication middleware. Due to various connectivity limitations, certain recommender systems (e.g., those running on personal and mobile devices) may be partially available online or their communication throughput may be limited. In some cases, the mediation may require time consuming data processing by one of the systems, which may prevent provision of real-time personalization services.

The fourth challenge refers to the structural heterogeneity and incompleteness of the user modeling data. As mentioned earlier, current recommender systems usually refer to specific application domains, and their services are provided using specific recommendation techniques, which imply specific UM representations. The lack of a standard representation for the UMs and the specific storage and access requirements imposed by various recommendation techniques result in a situation where various systems collect user modeling data in different ad-hoc forms. This heterogeneity causes several problems: various techniques may store the preferences of the same user in different forms, the information in various systems may be conflicting or outdated, and may be influenced by various cross-lingual and cross-cultural dependencies, and so on. All these heterogeneities aggravate the mediation task, since it must support not only the integration of user modeling data, but also the resolution of inconsistencies and conflicts among the data obtained from various systems (Francisco-Revilla and Shipman 2004).

This work focuses on resolving only the latter challenge—the heterogeneity of the available user modeling data. Although this work stresses the importance of data sharing, privacy, and technical challenges in UM mediation, they fall outside its scope. However, the reader is referred to Berkovsky et al. (2007c) for a discussion of privacy-preserving data exchange methods in collaborative filtering recommender systems and to Rabanser and Ricci (2005) for a discussion of integrated business models in E-Tourism recommender systems.

This work presents and details a generic framework for the UM mediation. It starts with a generic data representation model, allowing various aspects of the recommendation process to be expressed: users, items, contexts, and the evaluation of the recommendation. Then, it describes several specific instantiations of the model, demonstrating its compatibility with some state-of-the-art recommendation techniques. It proceeds with a discussion of one of the main problems of recommender systems: the sparseness of the UMs, which foils provision of high-quality recommendations to the users. The mediation of UMs (and other user modeling data) is proposed as a solution to the sparsity problem. Then, the UM mediation process is comprehensively discussed. Here, the general steps of the mediation are described and four major types of mediation that can potentially be applied are defined: cross-representation (including two subtypes), cross-user, cross-item, and cross-context mediations. Then, this work shows practical examples and the results of experimental evaluations of three mediation mechanisms: one subtype of cross-representation mediation (Berkovsky et al. 2006a), the generalized variant of cross-item mediation (Berkovsky et al. 2007a), and cross-context mediation (Adomavicius et al. 2005). The results of the experimental evaluations demonstrate the potential benefits of the mediation of user modeling data, as in certain conditions (will be detailed in the relevant sections) it allows improving the quality of the personalized recommendations provided to the users. This work continues with a survey of related papers and, finally, it presents current conclusions that can be drawn, and discusses several directions of future UM mediation research.

2 Data representation model

A unified model for user modeling data representation is required in order to facilitate the UM mediation. The discussion starts with a two-dimensional representation which is an abstraction of the data representation adopted by most of the existing systems. This model is then extended to a three-dimensional representation reflecting the context-awareness aspects.

2.1 Two-dimensional representation of user models

Most of nowadays recommender systems base the warehousing, i.e., storage, access and retrieval of their UMs, on a two-dimensional matrix representation. The two *generalized* dimensions of this representation are the *users* and the *items*. These dimensions are referred to as generalized because they may be described by sets of specific features. Hence, if the user is described by n features and the item by m features, the space of all possible user and item pairs is described by an $n + m$ dimensional space. For instance, when the users and the items are described by their unique identities only, the space of all possible users and items pairs is two-dimensional, where the first dimension refers to the user identifier and the second to the item identifier. In this case, the ratings given by users to items are described by a map from the two-dimensional space to a numeric range, e.g., $\{1, 2, 3, 4, 5\}$. In a more concrete way, in this situation, an $N \times M$ matrix (representing N users and M items) either represents directly or reflects the ratings given by the users to the items. The ratings are given on a predefined scale and can be given in an explicit or implicit way. Explicit ratings are typically provided by the users, while implicit ratings are inferred by the system through observing user behavior indicators. For example, if the user bought the recommended product, the system implicitly interprets it as a positive rating.

When the users and the items are described by sets of features, the matrix is still referred to as the description of ratings given by users to items. However, such a matrix is a high dimensional matrix, i.e., it is a function R'_{gen} from the $n + m$ dimensional space of $User_{feat} \times Item_{feat}$ to a set of ratings:

$$R'_{gen}\colon User_{feat} \times Item_{feat} \rightarrow rating.$$

In the above definition $User_{feat}$ represents the user features , $Item_{feat}$ represents the item features, and *rating* represents the ratings given by the users to the items.

In fact, this function is not defined for all the possible user and item pairs, i.e., the system may not know the rating values given by a user to all the items. Hence, given a user who requests a recommendation, the goal of a recommender system is: (1) to estimate the rating value for some items, which the user has not previously rated, and (2) to suggest some items to the user, for instance those items having maximal predicted rating. Note that the actual recommendation task heavily depends on the exact functionality (and service) provided by the system. This can include recommending the best item, ranking N best items, filtering highly irrelevant items and many others (Adomavicius and Tuzhilin 2005a). In the rest of this work the recommendation task refers to predicting a future rating, assigned by a certain user u to a certain item i.

Although R'_{gen} was defined as a function from a two-dimensional matrix, both of its basic dimensions $User_{feat}$ and $Item_{feat}$ can be described using a multidimensional representation by a set of features. However, since the described data representation is not based on any commonly agreed ontology, the separation between the basic dimensions of $User_{feat}$ and $Item_{feat}$ is ambiguous and somewhat artificial. Some systems may classify certain ephemeral features as features describing the users, while other systems may classify them as features describing the items. For example, consider a travel recommender system and a feature representing the *season* of the travel. This feature could be interchangeably considered as a one of the destination features (e.g., St. Moritz in winter) or as one of the user features (e.g., a user searching for a holiday resort in winter). To overcome this, the whole representation can be considered as a single multi-dimensional space of features, which reflects a single integrated list of features, where certain sets of features can be grouped into the basic dimensions of $User_{feat}$ and $Item_{feat}$.

This R'_{gen} representation is applicable to a wide variety of state-of-the-art recommendation techniques, as can be seen in the following examples:

- In collaborative filtering (Herlocker et al. 1999), the two-dimensional matrix R_{CF} is referred to as the *ratings matrix* and is represented by:

$$R_{CF}\colon User_{id} \times Item_{id} \rightarrow rating,$$

where $User_{id}$ and $Item_{id}$ are the unique one-dimensional identifiers of users and items,[1] and *rating* is the rating given by the user to the item. In this case, an individual UM is represented by a set of ratings given by the relevant user, and is referred to in the literature as the *ratings vector*. For example, consider the following ratings matrix R_{CF} = {((*Alice, "The Lord of The Rings"*),*1*), ((*Alice, "The Matrix"*), *0.8*), ((*Bob, "Psycho"*), *0.2*), ((*Bob, "Friday the 13th"*), *0*)}, representing the movie ratings of two users, Alice and Bob, given on a continuous scale of ratings between *0* and *1*. Typically, collaborative filtering systems do not store any additional information about the features and content characteristics of users and items, besides their identities.

- In content-based filtering (Morita and Shinoda 1994), the two-dimensional matrix R_{CB} is represented by:

$$R_{CB}\colon User_{id} \times Item_{feat} \rightarrow rating,$$

where $User_{id}$ represents the unique identifier of the users, $Item_{feat}$ represents a feature space describing the item's features, and *rating* reflects the user's ratings (e.g., in form of weights) to the items characterized by that feature. In this case, the UM is represented by the values reflecting the ratings given by $User_{id}$ to certain $Item_{feat}$ features, originating in the descriptions of items. For example, content-based matrix R_{CB} from the previous example can be $R_{CB} = \{((Alice, science\text{-}fiction), 0.9), ((Bob, horror), 0.1)\}$. In this example, each movie is described by a content-related feature representing the genre of the movie, taking in these examples the values *horror* and *science-fiction*. It can be observed that in content-based recommender systems the *raw* UMs contain the ratings of the users to the items described by a set of features. This information is typically used to build a refined UM that depends on the specific classification technique used by the recommender system.

- In demographic filtering (Krulwich 1997), the two-dimensional matrix R_{dem} is represented by:

$$R_{dem}\colon User_{feat} \times Item_{feat} \rightarrow rating,$$

where $User_{feat}$ represents a set of features describing certain demographic characteristics of a group of users to which the user belongs, $Item_{feat}$ represents either the unique identity of the item or a set of features reflecting item's content, and *rating* virtually represents the ratings given by a group of users with certain demographic characteristics to the items. In this case, the UM is represented by a combination ($user_{feat}$, $item_{feat}$) and the ratings provided by the user described by $User_{feat}$ to the items, containing $Item_{feat}$. For example, R_{dem} of the users Alice and Bob from previous examples can be R_{dem} = {(*female*, *science-fiction*), 0.9), ((*male*, *horror*), 0.1)}. In this example it should be pointed out that the very notion of user and item can depend on the specific recommendation technique. Here, for instance, the users are

[1] Although in this case the identity of the user (item) is considered as one of the features, it should be stressed the identity is a special feature, which facilitates a unique identification of user (item) in all the systems.

represented by a single user stereotype defined only by the gender feature. Similarly, the items are described only by the genre feature.

All the previous recommendation techniques could, in principle, adopt the generalized user and item representations, as in the demographic filtering approach, which generalizes the user description, and content-based approach, which generalizes the item description. In fact, many recommender systems providing personalized recommendations based only on the ephemeral session data, adapt the generalized user and item representation to the specific needs of the systems and exploit this representation for the purpose of generating the recommendations (Ricci et al. 2006).

In addition, also the way the ratings of the users are represented is highly heterogeneous across various recommender systems. Some systems store numeric ratings given by the users on a predefined, but not standard, scale (e.g., the scale may be discrete or continuous, the range of possible values may vary from one system to another, and so on), some store symbolic ratings (e.g., positive or negative ratings, thumbs-up or thumbs-down, and so on), some store system-specific feedback derived from user behavior (e.g., examining or not the recommended item, purchasing or not the recommended product, and so on), some store the resultant navigation history of the user (e.g., opening or not the recommended Web-link, period of time spent viewing the recommended Web-page, and so on), and others store the free-text feedback provided by the users (Hanani et al. 2001).

Several types of the user feedback are discussed and compared in Montaner et al. (2003). Such feedback is classified to four categories:

- No feedback—modification of the user data is done manually by the users using a specific component provided by the system.
- Explicit feedback—explicit opinion provided by the user. For example, numeric ratings assigned to artists or music bands (Shardanand and Maes 1995), annotations to viewed documents (Goldberg et al. 1992), or binary opinions regarding the interest value of Web-pages (Pazzani and Billsus 1997).
- Implicit feedback—user opinion is inferred by the system from monitoring user's behavior. For example, analyzing lists of preferred leisure activities (Krulwich 1997), reading times of the received messages (Morita and Shinoda 1994), or usage data of hypermedia systems (Kobsa et al. 2001).
- Hybrid feedback—combines both the explicit and implicit user feedback, such as in Joachims et al. (1997), Resnick et al. (1994), and Sakagami and Kamba (1997).

To resolve this heterogeneity and to refer to the wide variety of user feedback to the provided recommendations in a uniform manner, all the possible types of feedback are generalized and denoted by the term *evaluation* (McNee et al. 2003).

In order to address the heterogeneity of the evaluations assigned by (or predicted for) the user to an item, the R'_{gen} function was generalized to *experience* of a user for an item. An experience is defined as an evaluation function that maps a pair, the user that had the experience and the item experienced by the user, to an evaluation. An experience evaluation details how the user and the item are linked together. Formally, the experience is represented by:

$$Exp : User_{feat} \times Item_{feat} \rightarrow evaluation$$

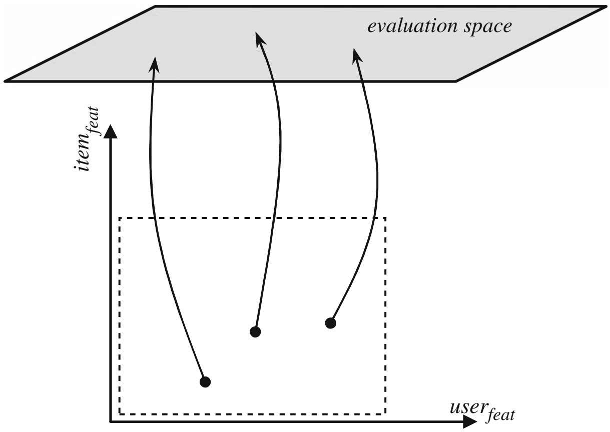

Fig. 1 Representation of experiences and their evaluations in two-dimensional space

where $User_{feat}$ and $Item_{feat}$ represent the feature spaces describing user and item features, and *evaluation* represents the feedback given by the user described by $User_{feat}$ for the item described by $Item_{feat}$. Figure 1 schematically illustrates the representation of experiences in a two-dimensional space.

For example, consider an experience *e* described verbally by "Alice likes science-fiction movies." Using a simple object-oriented-like notation, this experience can be represented by *Exp*(*user.name=Alice, item.movie.genre=science-fiction*)=*like*. This representation of the experience shows that the *evaluation* of a user, whose feature *name* is assigned the value Alice for the *item movie*, whose feature *genre* is assigned the value *science-fiction* is *like*.

This allows further generalization of the R'_{gen} representation of the matrix to the *Exp* representation for the UM warehousing comprising user- and item-dependent representation of experiences and evaluations. Also, over the *Exp* representation, a single UM, i.e., the model of a concrete user, is considered as a set of the *Exp* contents restricted to values of the features of this user. In other words, the user model for user *u* is the range of the *Exp* function restricted to the user *u*. Moreover, the *Exp* representation of experiences allows devising the following formulation: "*the recommendation is a task of predicting future evaluation of a new experience for a specific combination of* ($user_{feat}$, $item_{feat}$) *values, based on a set of past experiences*." In other words, the recommendations are aimed at predicting evaluations of the new experiences using the knowledge obtained from past experiences.

2.2 Three-dimensional representation of user models

The above generalization of the classical recommendation problem does not overcome a severe limitation of the majority of recommendation techniques: ignoring the context of the experience (Buriano et al. 2006). There is a variety of definitions for the term *context* in the literature. For example, it can refer to the user location, the time or the temperature of the day (Brown et al. 1997), or it can be considered as the subset of physical and conceptual states of interest to a particular entity (Pascoe 1998). One of the most comprehensive definitions of context is given in Dey et al. (1999): "*Context is any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves*."

With respect to recommender systems, Goker and Myrhaug (2002) defines context as a description of aspects of a situation and splits the generic user context into five components: (1) environment context—captures the entities that surround the user; (2) personal context—captures the state of the user and consists of two sub-components, the physiological context and the mental context; (3) task context—captures what the persons (actors) are doing in this user context; (4) social context—captures the social aspects of the current user context, such as friends, enemies, neighbors, co-workers and so on; and (5) spatiotemporal context—captures aspects of the user context relating to the time and spatial extent for the user context.

In fact, user preferences represented by the UM are generally valid only within specific contextual conditions, such as spatial, temporal, emotional, and other conditions. That is, a user's preferences stored in the UM may change as a function of various contextual conditions. Nonetheless, the generalized matrix representation *Exp* considers an experience as a user- and item-dependent entity only, not influenced by the contextual conditions, which may actually affect the evaluation of the experience. For example, two experience evaluations may be defined as "*Alice likes to see comedy movies with her friends*" and "*Alice does not like to see comedy movies with her parents*." In this example, if the companion of a user is treated as a contextual condition, the evaluation of the same experience is positive in one contextual condition and negative in another. Hence, to facilitate provision of context-aware recommendations, the above two-dimensional (user- and item-dependent) representation *Exp* should be extended by a third general dimension, reflecting various contextual conditions and features that may be considered by the recommender system.

The context-awareness issue has lead to a multidimensional warehousing of the UMs that captures the dependencies between the ratings and a generalized user-, item- and context-dependent model (Adomavicius et al. 2005). This model extends the two-variables function *Exp*, ignoring the context-awareness issue, to a three-variables function Exp_{CA}, incorporating a third dimension of context. Given the above generalization of ratings to the experiences *evaluation*, context-aware experience is defined by:

$$Exp_{CA}\colon User_{feat} \times Item_{feat} \times Context_{feat} \rightarrow evaluation.$$

Figure 2 schematically illustrates the representation of context-aware experiences in a three-dimensional space.

This representation, in addition to the standard $User_{feat}$ and $Item_{feat}$ features, also includes $Context_{feat}$ that represents the contextual conditions (or the values of the contextual features) of the experience. Similarly to $User_{feat}$ and $Item_{feat}$, $Context_{feat}$ is also described using a multidimensional representation by a set of features. Hence, a specific contextual condition of the experiences is referred to as a subspace of this multidimensional contextual space. For instance, in the above mentioned example,

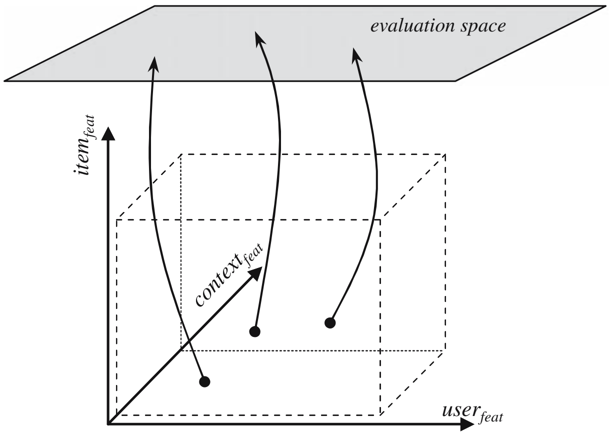

Fig. 2 Representation of context-aware experiences in three-dimensional space

only one contextual feature of *companion* out of a large $Context_{feat}$ set of contextual features is mentioned. When the *companion* feature is assigned the value of *friends*, the evaluation is positive, and when it has the value *family*, the evaluation is negative.

It should be stressed that modifying the representation of *Exp* function to Exp_{CA}, i.e., incorporating the contextual features, does not have any effect on the UM warehousing. UMs are still referred to as collections of past experiences, whereas the experiences are now context-aware. Also, although the definition and representation of experience was modified to capture the context-awareness issue, the definition of the recommendation task remains unchanged. This still consists of predicting the evaluations of the new experiences based on a set of past experiences and their evaluations. However, since the experiences were modified to include the contextual features $context_{feat}$, the recommendations should be provided in a context-aware manner, i.e., they should refer to a certain combination of $user_{feat}$, $item_{feat}$ and $context_{feat}$ values, and not of $user_{feat}$ and $item_{feat}$ values only.

Previous observation regarding the ambiguous separation of features between the basic dimensions is still valid in the three-dimensional representation. Various recommender systems can misinterpret the same feature and classify it to different basic dimensions. For example, the above feature of the *season* in a tourism recommender system in the previous section can naturally be considered as a contextual feature. Hence, although the third basic dimension of context was introduced, the whole representation can still be considered as a single multi-dimensional space of features, where certain sets of features can be grouped into the basic dimensions of $User_{feat}$, $Item_{feat}$, and $Context_{feat}$.

3 Mediation of user models for context-aware recommendations

The main problem in providing high quality context-aware recommendations using the context-aware Exp_{CA} representation for the UM warehousing is the sparseness of the data stored in the UMs, i.e., the lack of sufficient user modeling data about the user in specific contextual conditions (Ricci 2006). The problem of insufficient information for generating high-quality recommendations is a well-known problem of traditional recommender systems (Linden et al. 2003), which rely only on the two-dimensional representation R'_{gen} and ignore the contextual dimensions. This problem worsens when the contextual information is considered, as the initially sparse two-dimensional experiences are partitioned among multiple contexts, reflecting the specific contextual conditions of the experiences. As a result, the amount of available user modeling data referring to a specific contextual condition significantly decreases when the context-awareness issue is taken into account. Hence, a major question refers to tradeoff between the specialization of context-aware recommendation generations and the reduction of the available user modeling data.

This work discusses an approach aimed at overcoming the sparseness problem using a *mediation* of UMs and user modeling data. The exact definition of mediation is formulated as follows: "*mediation of UMs is a process of importing the user modeling data collected by other (remote) recommender systems, integrating them and generating an integrated user model for a specific goal within a specific context*." In this definition, the term integration refers to a set of techniques aimed at resolving the heterogeneities and inconsistencies in the obtained data. The mediation process facilitates instantiating the UMs through inferring the required user modeling data from past experiences and their evaluations in a three-dimensional context-aware representation space. Hence, it enriches the existing UMs (or bootstraps empty UMs) in the target recommender system using the data collected by the remote systems, and facilitates provision of better context-aware recommendations. In the rest of this section a general architecture of UM mediation will be presented and four practical mediation methods will be extensively discussed.

3.1 User modeling data mediation architecture

Two parties are involved in the mediation process. On the one hand, there is a *target* recommender system, i.e., the system requested to provide personalized recommendations to the user. Formally, this system acts as the initiator of the mediation process by requesting the available user modeling data from other systems. On the other hand, there are numerous *remote* recommender systems that may provide relevant user modeling data (i.e., past experiences) to the target recommender system. More precisely, these might not be recommender systems only, but also various services, Websites, sensors and even user's personal devices, that collected past experiences of the user. These two parties are interconnected via the UM mediator, which constitutes the core element of the mediation process. The general architecture of the UM mediation process is illustrated in Fig. 3.

As discussed earlier, the main difficulty of the UM mediation and the main focus of this work is overcoming the heterogeneity of the user modeling data. For example, recommender systems from different application domains imply different user modeling data stored in the UMs. Even within the same domain, different systems may store different information in their partial UMs, according to the specific recommendation technique being exploited (e.g., ratings vector in collaborative filtering UMs

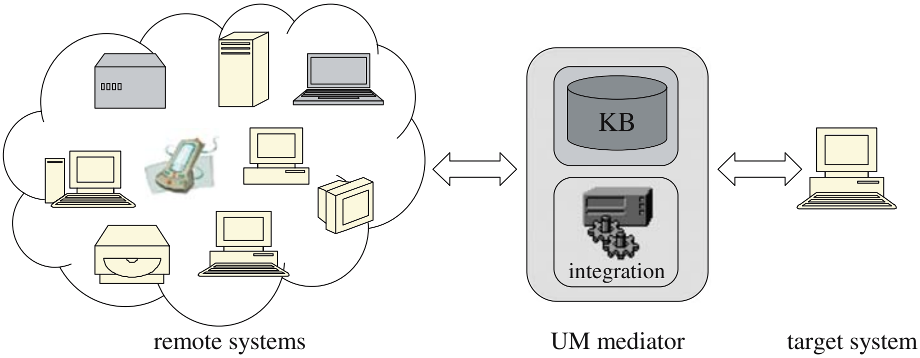

Fig. 3 Architecture of the user model mediation

(Herlocker et al. 1999) versus a feature vector of interest topics in content-based UMs (Morita and Shinota 1994)). Moreover, even the UMs of two recommender systems from the same application domain exploiting the same recommendation technique may use different terms to describe equivalent underlying objects, i.e., users, items, or domain features.

Hence, successful completion of the UM mediation task requires: (1) developing and applying reasoning and inference mechanisms for converting user modeling data between various representations, applications and domains, and (2) identifying and exploiting semantically enhanced knowledge bases, actually facilitating the above reasoning and inference. Combinations of various reasoning and semantic tools will allow identification of commonalities between the user modeling data representation of various systems. As a result, the mediator consists of two principal components:

- *Integration Mechanism.* The obtained past experiences may be represented in different ways, e.g., using various ontologies, domain-specific and application-specific terminology, or even in different languages. In addition, the evaluations of the same experience in different systems may be contradictory. Hence, this component is responsible for resolving conflicts and heterogeneities in the obtained user modeling data using various reasoning and inference mechanisms. This implies the integration mechanism to implement and apply certain policies for conflicts resolution in the obtained data.[2]
- *Knowledge Base (KB).* This is an auxiliary component, used by the integration mechanism. It contains semantically enhanced inter-domain and intra-domain knowledge bases representing dependencies and relationships between various user, item and context features. The data stored in the knowledge bases facilitate resolving the heterogeneities in the obtained user modeling data. For example, it allows reconciliation of the ontologies exploited by various recommender systems, converting the terms used by certain systems to a standard representation, and even provides machine translation tools resolving cross-lingual dependencies.

[2] It is reasonable to assume that various systems will provide user modeling data with different levels of accuracy and up-to-date information. Although the paper highlights the importance of resolving such conflicts, developing conflict resolution policies falls outside its scope.

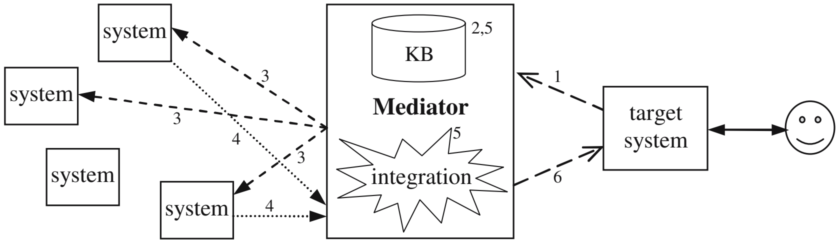

Fig. 4 Stages of the UM mediation

The envisioned flow of the user modeling data mediation process consists of the following stages (as illustrated in Fig. 4):

1. The recommendation request is treated by the target system as a request for a prediction of the evaluation of the new experience for a specific combination of $user_{feat}$, $item_{feat}$, and $context_{feat}$. By predicting this evaluation, the target system can also determine whether the item should be recommended to the user in a given context. The target recommender system queries the mediator for the UMs, containing past experiences that are relevant for predicting the evaluation of the new experience. The query contains the required ($user_{feat}$, $item_{feat}$, $context_{feat}$) combination.
2. The mediator analyzes the query and determines the set of remote recommender systems, which may store potentially relevant past experiences. This analysis can incorporate, for instance, external semantic data provided by the knowledge base, the distances between the domains, the representation technologies used in the systems, the availability of the systems, level of details available for the specific users, and many other factors.
3. The mediator forwards the query to the set of remote recommender systems that was determined in the previous stage.
4. Remote recommender systems, which actually store the relevant experiences, respond to the query and send to the mediator their locally collected UMs and/or the relevant experiences only.
5. The mediator integrates the obtained experiences using the semantic data provided by the knowledge base. Clearly, different combinations of UM representations in remote and target systems will require different integration mechanisms.
6. The generated user modeling data are sent to the target recommender system. Since the user modeling data of the target system are enriched in comparison to the locally collected data stored before the mediation, the system is capable of providing better recommendations to the user.

To illustrate the mediation flow, consider again the above example of a network of recommender systems dedicated to digital entertainment. The network consists of music, movies, TV programs, books, and humor recommender systems. Consider Alice, one of the users of the movies recommender system who requests the system to suggest a movie. To provide a better recommendation, the target movies recommender system queries the mediator for the relevant user modeling data (step 1). The mediator analyzes the request and the data provided by the movie recommender system (e.g., past opinions of Alice on the movies she has already seen and a list of potentially recommendable movies in theaters tonight) and identifies the set of remote recommender systems that can potentially provide relevant past experiences (step 2). Imagine that these systems are TV, books, music, and jokes recommender system. The mediator forwards the query for the relevant past experiences to these systems (step 3).

The remote recommender systems, which store the relevant experiences, send them back to the mediator (step 4). Imagine that only TV programs and books recommender system stored the relevant experiences: the TV programs system sent the list of programs seen by Alice during the last week and the books system send the list of books purchased by Alice through the Web-site of the books recommender system. The mediator integrates the acquired past experiences into a single UM using the knowledge base and converts it to the format required by the specific recommendation technique exploited by the target movies recommender system (step 5). For example, it mines the information available about the TV programs seen by Alice, extracts the topics of these programs, and checks whether there are recommendable movies with overlapping or similar topics. The mediator also identifies the writers of the books purchased by Alice and checks whether there are recommendable movies based on the novels written by these authors. Finally, the derived user modeling data (i.e., the list of topics and authors) is forwarded to the movies recommender system (step 6), where it is used for generating personalized recommendations.

3.2 User model mediation methods

The above scenario immediately raises a question: "What user modeling data are relevant for the mediation and need to be imported from the remote systems"? In principle, any available user modeling data (i.e., any past experience) may be relevant to some extent as input to the mediation process, since they may help predict the evaluations of the new experience. For example, consider a target recommender system that is supposed to predict the new experience evaluation for a specific combination of ($user_{feat}$, $item_{feat}$, $context_{feat}$) values. The possible groups of ($user_{feat}$, $item_{feat}$, $context_{feat}$) combinations in past experiences stored by other recommender systems are as follows (the respective mediation methods will be extensively discussed later in this section):

- Experiences that refer to the same combination ($user_{feat}$, $item_{feat}$, $context_{feat}$). They represent past experiences of the same $user_{feat}$ for the same $item_{feat}$ in the same $context_{feat}$, where the values of certain experience features, referring to the same objects, may be represented in different ways.
- Experiences where the values of two features are the same, and the value of one feature differs. Three possible combinations are:
  - ($user'_{feat}$, $item_{feat}$, $context_{feat}$)—past experiences of another $user_{feat}$ for the same $item_{feat}$ in the same $context_{feat}$.
  - ($user_{feat}$, $item'_{feat}$, $context_{feat}$)—past experiences of the same $user_{feat}$ for another $item_{feat}$ in the same $context_{feat}$.
  - ($user_{feat}$, $item_{feat}$, $context'_{feat}$)—past experiences of the same $user_{feat}$ for the same $item_{feat}$ in another $context_{feat}$.

- Experiences where the value of one feature is the same, and the values of two features differ. Three possible combinations are:
  - $(user_{feat}, item'_{feat}, context'_{feat})$—past experiences of the same $user_{feat}$ for another $item_{feat}$ in another $context_{feat}$.
  - $(user'_{feat}, item_{feat}, context'_{feat})$—past experiences of another $user_{feat}$ for the same $item_{feat}$ in another $context_{feat}$.
  - $(user'_{feat}, item'_{feat}, context_{feat})$—past experiences of another $user_{feat}$ for another $item_{feat}$ in the same $context_{feat}$.
- Experiences where the values of all three features are different. These experiences refer to $(user'_{itfeat}, item'_{feat}, context'_{feat})$ and represent past experiences of another $user_{feat}$ for another $item_{feat}$ in another $context_{feat}$.

Clearly, the first group of experiences is the most important for UM mediation, as it provides past evaluations of the target user for the required item in the relevant context. Such experiences require integration mechanisms for resolving possible heterogeneities in the representations of feature values or experience evaluations to be applied. The second group of experiences (with one feature different from the required combination) is also important for the mediation. These experiences represent past evaluations, where the values of two out of three features match the values of the features in the new experience. In this case, the mediation requires inference mechanisms to be applied to identify the relationships between the different values of the feature that differs, and to project the available evaluations onto the new experience.

In the third group of experiences, the values of two out of three features are different, and only one feature matches the values of a feature in the new experience. Hence, mediation of such experiences requires applying more complicated inference mechanisms (e.g., several consecutive inferences similar to the inferences from the previous group of experiences, where the value of only one of the features was different). Although this user modeling information may be relevant and may enrich the user modeling data in the target system, applying complicated inference mechanisms may 'deteriorate' the original data represented by the past experiences. Therefore, such experiences are currently not used in the mediation process. Obviously, the situation is even worse for the fourth group of experiences, where the values of all three features are different, and the mediation requires three inference mechanisms to be applied. Hence, these experiences are also not used in the mediation process.

In summary, two groups of experiences that may be considered as input for the mediation process are: (1) experiences having the required values of all three features, and (2) experiences having the required values of two features and a different value of one feature. The following analysis yields four particular types of UM mediation over the context-aware three-dimensional representation of experiences.

The first type of mediation is conducted between experiences having the required values of all three features, i.e., between heterogeneous representations of the same experience. Such mediation is referred to as *cross-representation mediation*. The other three mediation types are referred to as *cross-dimension mediations*. They are conducted over the experiences having the required values of two features and a different value of one feature. This means that the values of two out of three dimensions in the space are fixed and the mediation is performed across the third dimension.

Three types of cross-dimension mediations are possible: (1) *cross-user mediation*, where the values of item and context features are fixed and the user in the experiences is allowed to be modified; (2) *cross-item mediation*, where the values of user and context features are fixed and the item in the experiences is allowed to be modified; and (3) *cross-context mediation*, where the values of user and item features are fixed and the context in the experiences is allowed to be modified.

Figure 5 schematically illustrates the three cross-dimension mediations in a three-dimensional space. The top-left chart represents only the new experience, i.e., a certain combination of ($user_{feat}$, $item_{feat}$, $context_{feat}$) features, where the user's future evaluation needs to be predicted. Three other charts represent past experiences that can be imported at various types of cross-dimension mediation: the top-right chart represents the experiences imported at cross-user mediation, the bottom-left represents the experiences imported at cross-item mediation, and the bottom-right represents the experiences imported at cross-context mediation. In all the charts, the black dot represents the new experience, where the evaluation needs to be predicted and the circles represent past experiences that are imported and integrated at the respective type of the mediation. Note that cross-representation mediation is not shown graphically in Fig. 5. This type of mediation can be considered as mediation conducted between experiences stored in a single cell of the three-dimensional space, i.e., between the experiences stored in the black dot representing the new experience. In the rest of this section, all four possible types of UM mediation are extensively discussed.

3.2.1 Cross-representation mediation

This mediation is aimed at resolving the heterogeneity in the representations of the experiences of the same user for the same item in the same context. In other words, it incorporates past experiences of the same user for the same item in the same context, but expressed in different ways. For example, consider the following representation of the same item, a movie "*Gone with the Wind*", in two datasets: EachMovie (McJones 1997) and MovieLens (Herlocker et al. 1999). In EachMovie, it is classified as a *classic* movie, while in MovieLens it is classified as a *drama, romance* and *war* movie. In addition, a movie evaluation in EachMovie is a number between 0 and 1, while in MovieLens it is expressed by a number of stars on a 5-star scale. To implement mediation between these two systems, the mediator should be able to cope with such heterogeneities.

Hence, cross-representation mediation can be considered as an integration of user modeling data between heterogeneous representations of the values of the experience features. This means that although different representations of the features refer to the same underlying objects, they are expressed in different ways. As a result, the mediation is conducted between past experiences, where the values of all the components $user_{feat}$, $item_{feat}$ and $context_{feat}$ imply the same data, but are represented differently. This type of mediation is divided into two groups:

- Different representation of $user_{feat}$, $item_{feat}$ and $context_{feat}$ values. This mediation deals with a situation where the representation of one (or several) of the experience components is heterogeneous. This means that although the $user_{feat}$, $item_{feat}$ and $context_{feat}$ are semantically identical and reflect the same user modeling data, one

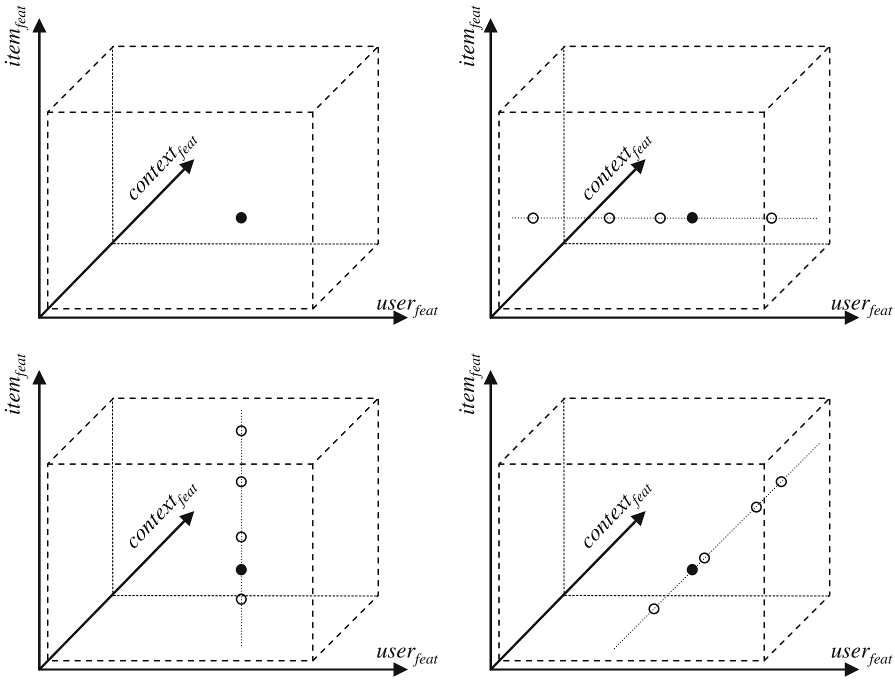

Fig. 5 Cross-dimension mediations: top-left—the new experience to be predicted, top-right—experiences imported at cross-user mediation, bottom-left—experiences imported at cross-item mediation, bottom-right—experiences imported at cross-context mediation

(or several) of them is (are) syntactically expressed in different ways. For example, collaborative filtering systems represent an item using its unique identifier only, while content-based systems represent the item using the set of its features. This mediation requires inference mechanisms to be applied identifying commonalities and dependencies between various representations of semantically identical $user_{feat}$, $item_{feat}$, or $context_{feat}$, using an external domain-specific knowledge base. Hence, this variant of cross-representation mediation is referred to as *cross-technique mediation* as the difference in the representations reflects the differences in the recommendation techniques (Berkovsky et al. 2006a).

- Different representations of the evaluation values. In this mediation, the representations and the values of $user_{feat}$, $item_{feat}$ and $context_{feat}$ features are identical, but the evaluations of the experiences are expressed in different ways, i.e., the heterogeneity is in the representation of the evaluations. For example, the target recommender system represents the evaluation as a discrete numeric rating on a scale between 1 and 10. However, the remote system represents it as positive or negative evaluation only. To overcome this heterogeneity, it is necessary to map and reconcile the evaluation representation values between the two scales. This case is distinguished from the previous one since overcoming the heterogeneity of feature representations is considered a more complicated task than the mapping between the evaluation scales.

3.2.2 Cross-user mediation

Although the term mediation is not explicitly mentioned in collaborative filtering recommender systems, cross-user mediation and inference actually constitute the basis of this popular recommendation technology (Herlocker et al. 1999). Collaborative filtering is based on the assumption that people with similar tastes (i.e., people who agreed in the past) will prefer similar items (i.e., will agree in the future) (Shardanand and Maes 1995). Or, in a simplistic view, collaborative filtering recommends items liked by similar users.[3] In order to generate a recommendation, collaborative filtering systems initially create a neighborhood of users with the highest similarity to the active user, and then generate a recommendation by integrating the ratings of these users. Hence, this process can be considered as a cross-user inference, or a particular case of cross-user mediation, where the mediated user modeling data are the ratings of similar users. In addition, other variants of cross-user inference are applied in several existing recommendation techniques, e.g., demographic filtering (Krulwich 1997), collaborative by content recommendations (Pazzani 1999), and some hybrid approaches (Vozalis and Margaritis 2004).

It should be noted that the existing implementations of cross-user mediation in the state-of-the-art recommender systems mostly ignore the context-awareness issue. This means that they project the collected experiences onto the two-dimensional representation R'_{gen}, not reflecting the contextual conditions of the experience. Hence, these recommender systems apply inference mechanisms assuming that the collected experiences were recorded for the same contextual conditions. Thus, the prediction of the new experience evaluation can be considered as an inference process incorporating past experiences of other users for the same item in the same, actually *undefined*, context, i.e., the prediction generation process is pure cross-user inference.

3.2.3 Cross-item and cross-domain mediation

Cross-item mediation is also applied in various existing recommendation techniques, such as content-based filtering (Morita and Shinoda 1994), item-to-item collaborative filtering (Sarwar et al. 2001), utility-based recommendations (Manouselis and Sampson 2004), and in some hybrid approaches (Pazzani 1999). In general, these techniques assume that the similarity of items may also be used for providing personalized recommendations, i.e., items which are similar to the items the users liked in the past should be recommended to the users (most cross-user similarity metrics discussed in Herlocker et al. (1999) may be applied also for computing cross-item similarity).

In these systems, the context-awareness issue is also mostly ignored. The collected experiences are represented using the two-dimensional representation R'_{gen}, such that the collected experiences are considered as if they were all recorded in the same contextual conditions. Thus, the prediction of the new experience evaluation can be considered as the result of an inference process, incorporating past experiences of the

[3] The reader is referred to Herlocker et al. (1999) for a discussion on collaborative filtering similarity metrics.

same user for other items in the same, actually *undefined*, context. This means that the recommendation generation process is performed through cross-item inference.

To conduct cross-item mediation, the similarity of items (or, relationships between the items) should be defined for any arbitrary pair of items. However, not for any pair of items can the similarity be easily defined. For example, in item-to-item collaborative filtering (Sarwar et al. 2001), cross-item similarity is computed by means of user ratings to items. In this case, computation of the similarity between two items requires the items to be rated by a non-empty set of overlapping users. This requirement may be too strong for sparse ratings matrices. Moreover, in many conditions, the available past experiences do not necessarily reflect the user's evaluation for an individual item, but rather on a generalized group (or category) of items. For example, a user may express his opinion not on a single movie, but on a genre of movies, or on movies directed by a certain director.

Hence, in a broader view, the individual items need to be grouped. This allows a more complex type of mediation, incorporating the evaluations of past experiences for a generalized group of items, to be applied (Mehta et al. 2005). Generalizing individual items into groups and domains and then exploiting cross-domain dependencies and inferences introduces the issue of *cross-domain mediation*, where the evaluation of the new experience for a generalized group of items from a certain domain is inferred from past experiences for items in other domains (Berkovsky et al. 2007b). In this sense, cross-domain mediation can be considered as a mediation incorporating past experiences of the same user in the same contextual conditions, for another generalized group of items.

3.2.4 Cross-context mediation

The issue of cross-context mediation is a new research direction in user modeling. Such mediation is based on context-aware representation of the experiences, and its goal is to predict the evaluations of the new experiences in a given context using past experiences in other contextual conditions (Berkovsky et al. 2007a). This means that cross-context mediation incorporates past experiences of the same user for the same item in other contextual conditions. For example, it can predict future evaluation for an item by a user in the *evening* given past evaluation of the same user for the same item in the *morning*, or, it can predict future evaluation for an item by a user when accompanied by a group of friends given past evaluations of the same user for the same item when accompanied by a parent.

Since the state-of-the-art recommender systems mostly ignore the context-awareness aspect and are not capable of providing context-aware recommendations, this type of mediation requires the definition of various novel cross-context reasoning mechanisms. Two simple mechanisms, exploiting semantically enhanced OWL (Resource Description Framework) representations of $user_{feat}$, $item_{feat}$ and $context_{feat}$, were discussed in Berkovsky et al. (2007a):

- *Rule-Based Reasoning.* This reasoning mechanism exploits the semantically-enhanced representations of the experience components for the purposes of defining a set of reasoning rules that exploit the relationships between the values of the

features. For example, consider a semantic representation of times of day, and a reasoning rule defining that user's preferences regarding a certain item (e.g., stocks news report) in the *evening* are opposite to preferences in the *morning*. Or, consider another rule based on the same semantic representation of times of day, which defines a projection of user's preferences at 4PM onto a more general *afternoon* time period. Applying these rules facilitates the inference of the required user modeling data across various contextual conditions.
- *Similarity-Based Reasoning*. This reasoning mechanism exploits the semantically-enhanced representations of the experience components for the purposes of defining an explicit similarity metric, capable of computing the similarity between any arbitrary pair of contextual conditions. For example, such a metric may express similarity between Tuesday and Wednesday as mid-week days and dissimilarity between Tuesday and Sunday as mid-week and weekend days. Such a cross-context similarity metric allows the derivation of various adaptation rules, similar to the rules used in Case-Based Reasoning (Ricci et al. 2006), and facilitates reuse of the evaluations of past experiences. For example, this can be done by a collaborative-like weighted aggregation of the evaluations of past experiences of the same user for the same item in similar contextual conditions.

A comparison of the above discussed rule-based and similarity-based reasoning approaches shows that, on the one hand, rule-based reasoning may produce more accurate user modeling data, as the reasoning rules are typically defined by domain experts. On the other hand, defining and updating the inference rules in today's highly dynamic information world may hamper the scalability of the mediation process. Conversely, the typical scenario for similarity-based reasoning is fully autonomous and therefore gives a more flexible mediation process. However, similarity-based reasoning requires a large number of past experiences to bootstrap the reasoning process. Other machine learning approaches can be considered for this purpose.

4 User modeling mediation examples

To demonstrate and evaluate the benefits of UM mediation, two examples of practical mediation approaches were implemented and experimentally evaluated and a third example is presented by analyzing prior research: (1) *cross-technique mediation* (Berkovsky et al. 2006a), as a variant of cross-representation mediation, where the representation of the experiences is different due to the requirements imposed by the recommendation techniques, (2) *cross-domain mediation* (Berkovsky et al. 2007b), as a generalized form of cross-item mediation, where item-to-item relations are generalized to higher-level domain-to-domain relations, and (3) *cross-context mediation* as an analysis of prior research (Adomavicius et al. 2005).

The issue of cross-user mediation was studied intensively in many prior papers on collaborative filtering, such as Goldberg et al. (2001), Good et al. (1999), Herlocker et al. (1999), Pazzani (1999), Sarwar et al. (2001), Vozalis and Margaritis (2004) and many others. Collaborative filtering became one of the most widely used recommendation techniques researched and applied in a wide variety of domains and applications (Schafer et al. 2000). It has been evaluated many times and proved to be an effective and accurate recommendation technique (Herlocker et al. 2004). Hence, further evaluation of cross-user mediation through collaborative filtering seems to be redundant in the context of this work. For the exact experimental results and various analyses of collaborative filtering cross-user mediation, the reader is referred to the collaborative filtering papers mentioned above.

It is worth mentioning that one of the main challenges in conducting experimental evaluations of UM mediation is the lack of publicly available datasets, which represent the user modeling data of the same users across several applications and/or domains, or store heterogeneous representations of the experiences, or reflect the contextual conditions. Currently available datasets mostly contain only collaborative filtering data, representing users' ratings to items from three domains: movies (MovieLens (Herlocker et al. 1999) and EachMovie (McJones 1997)), books (BookCrossing (Ziegler et al. 2005)), and jokes (Jester (Goldberg et al. 2001)). However, these datasets are not cross-linked, i.e., no users from one dataset can be identified in another dataset and none of them contains contextual information about the ratings. For these reason, EachMovie dataset (McJones 1997) of movie ratings was used in the experimental evaluations of cross-technique and cross-domain mediation methods, whereas the evaluation of cross-context mediation required collecting a small dataset (Adomavicius et al. 2005) that reflects certain contextual conditions of the movie ratings. The details of the dataset preprocessing for the experiments will be provided in the relevant sub-sections.

All the mediations required exploiting semantically enhanced domain knowledge bases. In cross-technique mediation, semantic information about the movies was required for extracting and mining various features of the movies rated in the collaborative filtering UMs. In cross-domain mediation, semantic information was required for partitioning the movies according to genre and also for devising inter-domain distances (will be explained later). Conversely, in cross-context mediation, movies genre information was used as one of the features for generalizing the movie ratings provided by the collaborative filtering UMs to the genre ratings.

Since all the evaluated mediation methods used movie ratings, their implementations exploited semantic movie information provided by the IMDb, the Internet Movie Database. The IMDb provides information from 49 feature categories, such as *genre*, *actors*, *directors*, *writers*, *composers*, *keywords*, *distributors*, *languages*, *awards*, and many others. A small part of these feature categories was used in the implementations of the mediations (the details will be discussed later). To speed up the experimental evaluations, an offline version of the IMDb was downloaded, parsed, and inserted into MySQL database.

As no operational recommender systems were involved in the mediation, all the implementations were centralized. This means that the target recommender system, the remote system and the mediator were running on the same machine. The evaluations were programmed in Java language. In part of the implementations, the programming was done in a multi-threaded manner, to reflect the inherent distribution of the components and simulate a setting with several recommender systems. The details and algorithms of the above mediation methods, experimental evaluations, and their analysis are extensively discussed in the following sub-sections.

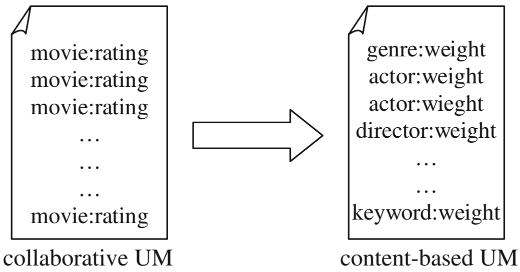

Fig. 6 Collaborative filtering to content-based mediation of user models

4.1 Cross-technique mediation

The first example demonstrates the possible benefits of cross-technique mediation. Cross-technique mediation deals with the heterogeneity of the experience representations in various recommendation techniques. As such it is considered as a variant of cross-representation mediation. In Berkovsky et al. (2006a), the mediation experiments focused on cross-technique mediation of user modeling data, where the experiences (and in particular, their components) were represented heterogeneously. The mediation was conducted between two recommender systems from the same application domain of movies. This mediation scenario assumed a user requesting recommendations from a content-based movies recommender system that has no prior knowledge about the user. However, it was assumed that there exists an external collaborative filtering recommender system that may provide some information about the preferences of this specific user.

The source of the user modeling data (i.e., the remote system) was a collaborative filtering (Herlocker et al. 1999) recommender system, whose user modeling information was mimicked by EachMovie dataset. In collaborative filtering, the UMs were represented by a list $UM_{CF} = \{i_1 : r_1, i_2 : r_2, \ldots, i_n : r_n\}$, where every pair $i_k : r_k$, corresponds to a rating r_k provided by the user to an item i_k. These UMs were mediated to a target content-based recommender system (Morita and Shinoda 1994), where the UMs were represented by a list $UM_{CB} = \{f_1 : w_1,\ f_2 : w_2,\ \ldots,\ f_n : w_n\}$, where f_k denotes one of the domain features and w_k denotes the level of the user's preference regarding this feature.

For example, consider the following collaborative filtering experiences of a user UM_{CF} = {"*The Lord of The Rings*":1, "*The Matrix*":0.8, "*Psycho*":0.2, "*Friday the 13th*":0, "Star Wars":0.9, "The Nightmare":0.1, "Alien":0.9}, where the evaluations are given on a continuous scale between 0 and 1. It can be inferred that the user likes science-fiction movies, and dislikes horror movies. Thus, the mediated content-based UM of the user may be UM_{CB} = {*science-fiction*:*0.9*, *horror*:*0.1*}, where the weights of the genres are computed as an average of the ratings given to the movies in this genre. Similarly, also the weights of other movie features, such as directors, actors and other can be computed. Figure 6 schematically illustrates the representation of the original collaborative filtering UM and the target content-based UM.

Obviously, applying such mediation requires external knowledge defining, for example, that "*The Matrix*" is a *science-fiction* movie and "*Psycho*" is a *horror* movie. Such movies metadata and content information were mined from the above mentioned IMDb database. Since the IMDb provides movie information from 49 feature categories, a wrapper feature selection algorithm (Kohavi and John 1997) was applied to identify the set of features that allows the highest accuracy of the generated recommendations to be achieved. As a result, only five feature categories were selected and used in the experiments: *genres*, *keywords*, *actors*, *actresses*, and *directors*.

This mediation converted the collaborative filtering UMs represented by the ratings vectors into content-based UMs represented by weighted vectors of features. The mediation was based on the underlying assumption that users' ratings implicitly reflect their preferences for the movie features, e.g., preferred movie genre, beloved topic, or favorite actors and directors. Hence, the mediation can be considered as a process of identifying and learning commonalities in the features of positively and negatively rated movies, and generalizing them into a weighted list of features.

For each movie rating in the collaborative filtering UM, a list of the movie features from the above five categories was extracted from the IMDb. Then, the weights of the movie features in the content-based UM were modified according to the rating assigned to the movie by the user. For example, consider a strongly positive rating of 0.8 that was assigned to a movie *"Star Wars"*. According to the IMDb, the genres of *"Star Wars"* are *action*, *adventure*, *fantasy* and *science-fiction*. Thus, the weights of these four features were increased by 0.8. Similarly, the weight of the movie director *George Lucas*, and the weights of all the actors involved in the movie and other movie features were increased by 0.8. To normalize the impact of the more frequent features, frequencies of the features were also computed. Hence, the frequencies of all the movie features were increased by 1.

When the above mediation was completed for all the rated movies in the collaborative filtering UM, the user was represented by the content-based UM as a set of feature weights $\{w_{i(1)}, w_{i(2)}, \ldots, w_{i(k)}\}$ for a subset of k features available in the above five feature categories, and the corresponding frequencies $\{c_{i(1)}, \ldots, c_{i(k)}\}$. The content-based UM facilitated generating content-based recommendations for movies. This was done by (1) extracting from the IMDb all the features of a recommendable movie from the selected five feature categories, and (2) computing the predicted rating as a weighted average of the weights of the features that appeared both in the UM and in the description of the movie. The predicted rating was computed by:

$$\text{rating}(m) = \frac{\sum_{i \in F(u) \cap F(m)} w_i c_i}{\sum_{i \in F(u) \cap F(m)} c_i}$$

where $F(u)$ denotes the set of features in the content-based UM and $F(m)$ denotes the set of the movie features. Given a set of recommendable movies, a movie with the highest predicted rating is recommended to the user.

The goal of the experimental evaluation was to determine the conditions in which cross-technique mediation improves the quality (in this case, measured in terms of accuracy) of the generated recommendations. To evaluate the effect of the mediation, the accuracy of the collaborative recommendations generated using the original collaborative filtering UMs was compared with the accuracy of the content-based recommendations generated using the mediated content-based UMs. In general, if the mediated UMs can generate recommendations, whose accuracy is comparable with

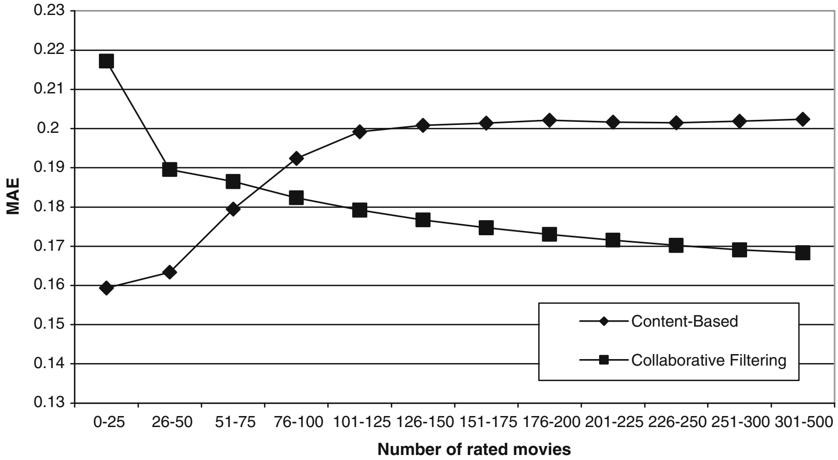

Fig. 7 Cross-technique mediation of user models

those obtained with the original data and the original technique, this could validate the feasibility of the proposed mediation approach. The more particular goal of the experiment was to assess the impact of the sparseness of the original collaborative filtering UMs on the accuracy of the generated recommendations.

For this, all the available users in the dataset were partitioned into 12 groups, according to the number of rated movies in their collaborative filtering UMs: below 25 movies, between 26 and 50 movies, and so on, up to users that rated between 301 and 500 movies. Collaborative filtering UMs were partitioned into 90% training set and 10% testing set. In other words, the training set of 90% of ratings served as a basis for the UM mediation, while the recommendations were generated on the ratings in the testing set. Two types of recommendations were generated: collaborative and content-based. Their accuracy was assessed using the widely-used MAE measure (Herlocker et al. 2004), reflecting the statistical accuracy of the generated recommendations:

$$\text{MAE} = \frac{\sum_{i=1}^{N} |p_i - r_i|}{N},$$

where N denotes the total number of generated recommendations, p_i is the predicted rating, and r_i is the real rating to the item number i in the testing set. Figure 7 shows the MAE values. The horizontal axis reflects the number of rated movies in the collaborative filtering UMs, while the vertical stands for the MAE values.

The results show that the MAE of the recommendations generated over content-based UMs, generated from collaborative UMs containing less than 50 ratings, is relatively low, around 0.16. This is explained by the observation that for a low number of rated movies in the collaborative filtering UM, it is easy to find the important features for content-based recommendation, i.e., the features that are strongly liked or disliked by the user, and to compute their weights accurately. For larger UMs, between 50 and 125 rated movies, the MAE linearly increases with the number of rated movies.

The probable reason for this is the larger number of features that appear in the rated movies. When the number of features in a content-based UM increases, it is harder to distinguish reliably between features that are important for accurate recommendations and features that introduce noise into the computation and hamper the accuracy of the generated recommendations. Finally, for the collaborative UMs with over 150 rated movies, the MAE stabilizes at around 0.20.

Comparison of the MAE values of content-based and collaborative recommendations shows that for groups of users that rated below 75 movies, content-based recommendations based on the mediated UMs significantly outperform collaborative recommendations based on the original UMs. To understand this result better, the distribution of users over the groups was computed (i.e., what percentage of users belongs to each one of the 12 groups of users). The distribution showed that 78.4% of the users in the dataset rated below 75 movies. Due to the sparsity problem (Linden et al. 2003), the accuracy of the collaborative filtering recommendations over their UMs is relatively low. Thus, improving the accuracy of the recommendations for these groups of users is extremely important and UM mediation from the collaborative filtering UMs to the content-based UMs and further generation of content-based recommendations provide a solid alternative technique.

For users with more than 75 rated movies in the collaborative filtering UMs, the accuracy of the collaborative filtering recommendations outperforms the accuracy of the content-based recommendations. It can be hypothesized that applying more accurate weighting mechanisms, e.g., weighting specific features within the feature categories or discovering dependencies between specific features, may improve the accuracy of the content-based recommendations also for larger UMs. In summary, the results show that in certain conditions (i.e., below 75 movies rated in the collaborative filtering UM) the above cross-technique mediation of UMs is beneficial, as it allows improving the accuracy of the recommendations, while in other conditions the available user modeling data are sufficient for generating accurate recommendations.

The implemented variant of cross-representation mediation mimicked a scenario, where the UMs of the target content-based recommender system were initially empty, and all the user modeling data were mediated from collaborative filtering recommender systems. Since the content-based UMs were initially empty, no content-based recommendation could be provided to the users, unless the UMs were generated by the mediation. Hence, this can be considered as an aggravated setting for the mediation, since the quality of the recommendations generated by the target system depended totally on the accuracy of the mediation. Other evaluations reported in the follow-up work, demonstrated that the mediation may be beneficial also in a setting, where a certain content-based user modeling data are available, as it still allowed the accuracy of the generated predictions to be improved.

This example showed the potential of cross-technique mediation in a specific scenario (use of collaborative UMs to bootstrap a content-based recommender system) and shed light on some practical characteristics of this scenario. One important observation regarding the above approach should be stressed. The content-based recommendations are generated solely based on content-based UMs, derived from collaborative filtering UMs. As such, the recommendation mechanism is capable of generating content-based recommendations regardless of the number of available movie ratings. Hence, this approach resolves the well-known *first-rater problem* (Morita and Shinoda 1994) of collaborative filtering, where an item cannot be recommended unless it has already been rated by a sufficient number of users. However, being a pure content-based recommendation mechanism, it suffers from the *serendipity problem* (Herlocker et al. 1999), i.e., it can recommend only movies that are similar to the movies that have already been rated by the user and cannot generate unexpected recommendations. Hence, while resolving one problem, the described approach introduces another.

4.2 Cross-domain mediation

The second example demonstrates the possible benefits of cross-domain mediation as a generalized form of cross-item mediation. In this case, the mediation scenario assumes that a user requested a recommendation from a domain-specific recommender system that has only partial or even no prior knowledge about the user. However, it was assumed that there exist external recommender systems for other domains that may provide some user modeling data about the preferences of the user. These data can be used by the target system for generating personalized recommendations despite the lack of domain-related UM. It should be noted that these data could be only partially related to the target domain data, and its utility could be limited. For example, if the target domain is movies and the external domain is CDs, the user modeling data regarding the user's CD preferences are not necessarily valuable for generating movie recommendations for the user.

We recall that cross-item mediation deals with importing and integrating past experiences of the same user for another item in the same context. However, such mediation requires the definition of item-to-item similarities or correlations, which in certain conditions, e.g., sparse collaborative filtering ratings matrix, cannot be defined for any arbitrary pair of items. To overcome this limitation, the items can be grouped according to a certain clustering criteria, such that the grouping facilitates computing group-to-group similarities, rather than item-to-item similarities. Such similarities allow a more complex type of mediation to be applied, importing and integrating past experiences for groups of items, rather than for individual items.

A simple criterion for generalizing items into groups is their domain classification, e.g., music, books, and so on. Generalizing individual items into domains introduces the issue of cross-domain mediation. It requires defining domain-to-domain similarities and providing exact inference mechanisms, capable of inferring the evaluation of the new experiences for a generalized group of items in a certain domain from the experiences for items in other domains. Hence, cross-domain mediation can be considered as a variant of cross-item mediation incorporating past experiences of the same user in the same contextual conditions, but for another group of items.

In cross-domain mediation, the user modeling data are mediated from remote recommender systems providing recommendations for items from other domains. For the sake of simplicity, cross-domain mediation in Berkovsky et al. (2007b) focused on a scenario where both the remote and the target systems exploited collaborative filtering recommendation technique (Herlocker et al. 1999). Hence, all the parties involved in the mediation represented their UMs using the ratings vector representation. For the cross-domain mediation case study, a centralized ratings matrix was virtually split into a set of domain-specific matrices according to the domain classification of the underlying items. This mimicked a distributed setting, such that every system from a certain domain stored its own (local) ratings matrix, whose set of items was restricted only to items that belong to the relevant domain. This split of the original centralized ratings matrix into a set of domain-related matrices can be considered as a vertical partitioning of the original matrix, as shown in Fig. 8. Note that according to this partitioning, domain-related sets of items are not disjoint, i.e., an item may belong to multiple domains. This setting is not uncommon in E-Commerce services (Schafer et al. 2000), where ambiguous categorization of items is explained by different classifications of products, their providers, or E-Commerce sites.

Over this domain-related partitioning, three types of user modeling data can be mediated, yielding three practical cross-domain mediation approaches. These types of data reflect the stages of the recommendation generation in collaborative filtering (Herlocker et al. 1999). For the *similarity computation* stage, the entire UMs (i.e., ratings vectors) can be mediated from the remote recommender systems to the target one. For the *neighborhood formation* stage, the list of nearest-neighbors, and their similarities computed by the remote systems, can be mediated. Finally, for the *recommendation generation* stage, complete recommendations generated by the remote systems can be mediated. It should be noted that all these mediation approaches should be compared with a baseline setting. It can be assumed that a setting where no user modeling data are being mediated and the recommendations are generated using only the data stored in the rating matrix of the target domain, is considered as the baseline setting for the comparisons. These recommendations are referred to in this work as *local* recommendations.

In the following, the mediation process and the generation of recommendations in each one of the above cross-domain mediation approaches will be briefly described:

a. The first mediation approach improves the traditional collaborative filtering by basing the *similarity computation* not only on the UMs available locally at the target system, but also on the UMs mediated from remote systems. To enrich the UMs of the target system, remote systems send to the target system their local UMs, i.e., their ratings vectors. Upon receiving the ratings vectors from the remote systems, the target recommender system constructs a unified ratings matrix by integrating the locally stored ratings vectors of the target system with the vectors mediated from the remote systems. Then, traditional collaborative filtering mechanisms of similarity computation, neighborhood formation and weighted recommendation generation are applied over the unified matrix. Since the unified ratings matrix can

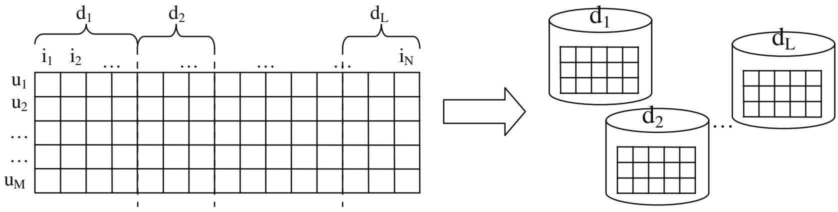

Fig. 8 Vertical domain-related partitioning of the ratings matrix

be considered equivalent to the centralized ratings matrix (if the UMs were not partitioned across domain-related matrices), this mediation approach is referred to as the *standard* approach.

b. The second mediation approach improves the accuracy of the *similarity computation*, by basing it not only on the similarity of the users computed by the target system, but also on their similarity values obtained from the remote systems (assuming that the users used both systems). Since the remote systems are collaborative filtering recommender systems, they can autonomously compute the similarity between the users using their local UMs. Hence, each remote system sends its list of most similar users and their similarity values to the target system. Upon receiving the sets of the most similar users, the target system computes the overall similarities between the target user and the other users as a weighted average of the similarities computed locally by the remote systems, such that the weights reflect the correlations between the target domain and the remote domains.[4] Then, a set of most similar users with the highest overall similarity values is selected, and standard collaborative filtering recommendations are generated. Since this approach incorporates similarity values computed by remote systems from various domains, it is referred to as the *cross-domain* approach.
c. The third mediation approach improves traditional collaborative filtering by *integrating recommendations* generated by the remote systems, rather then relying only on a single recommendation generated by the target system. Since in domain-related partitioning an item may belong to several application domains, any system generating recommendations for items from one of these domains can potentially generate the required recommendation. Hence, the remote recommender systems storing the ratings to the relevant item generate local collaborative filtering recommendations based on their local UMs, and send the recommendations to the target system. Upon receiving the recommendations, the target system integrates them into a single recommendation value by averaging the locally computed recommendation with the received recommendations. Since in this case, the mediation averages the values of the recommendations received from a set of target systems, this approach is referred to as the *remote-average* approach.

Experimental evaluation of these three mediation approaches and the baseline *local* approach was conducted using EachMovie dataset of movie ratings (McJones 1997). Since the ratings in EachMovie belong to a single domain of movies, the data were partitioned into domain-related datasets by separating the movies according to their genres, which were mined from the IMDb. As a movie typically belongs to several genres, such partitioning mimicked the above vertical domain-related partitioning of the centralized ratings matrix, where different domains may partially overlap. According to this partitioning, each domain was represented by the relevant genre ratings matrix and by the respective recommender system, generating collaborative filtering recommendations for movies from this genre.

[4] Inter-domain correlation between a pair of domains was computed as the average similarity of all the possible pairs of items that belong to these domains.

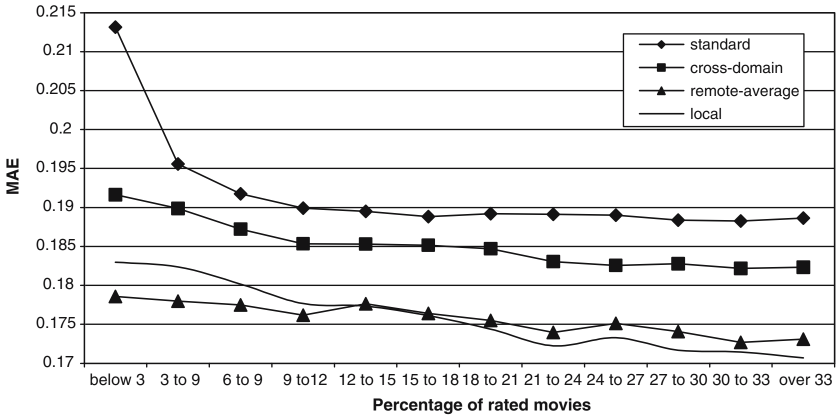

Fig. 9 Cross-domain mediation of collaborative filtering user models

The goal of this evaluation was to evaluate the impact of various mediation types as a function of the sparseness of the original UMs. Hence, the users in the genre-related datasets were partitioned into 12 groups, according to the percentage of movies rated in the relevant genre: below 3% of movies, between 3% and 6% movies, and so on, up to users who rated over 33% of the genre movies. For every group, 1,000 recommendations were generated for various combinations of the target user, predicted movie, and the target genre of the movie. Also in this experiment, the accuracy of the recommendations was assessed using the MAE measure (Herlocker et al. 2004). Figure 9 shows the results of the experiments. The horizontal axis reflects the percentage of rated movies in the target genre collaborative filtering UM, while the vertical axis stands for the MAE of various mediation approaches.

The results show that in most cases the *local* approach is still capable of generating accurate recommendations. Precisely, it outperforms the *standard* and *cross-domain* approaches for any percentage of movies rated in the target genre, while the *remote-average* approach outperforms it for a low percentage of rated movies. This means that the similarity of users (and the corresponding recommendations) in the *local* approach, which is computed using only the local data of the target domain, is more accurate than the similarities of users in the *standard* and *cross-domain* approaches, which are computed using also the ratings matrices of other domains.

Comparing the *local* and *remote-average* approaches shows that for a small percentage of rated movies, i.e., sparse ratings matrix, the *remote-average* approach is slightly more accurate (statistically insignificant). This can be explained by the observation that the predictions are generated using additional knowledge acquired by mediating data from other relevant domains, and not only using the data from the target domain. For a higher percentage of rated movies, the *local* approach is more accurate (also statistically insignificant), since the locally stored data is sufficient and the mediated data hampers the accuracy of the recommendations. Hence, the *local* recommendations outperform the *remote-average* ones. In conclusion, the accuracies of the *local* and *remote-average* approaches are very close.

However, it should be stressed comparing the performance of the mediation approaches using only the MAE is misleading, since in certain conditions the *local* and *remote-average* approaches may be inapplicable. For example, for a group of below 3% of rated movies, recommendations can be generated only for the domains with a large number of available movies, as for small domains, 3% of the number of movies may be lower than the minimal number of movies required for a reliable similarity computation. Moreover, a decent number of similar users is required for generating reliable recommendations. The *local* approach may not find a sufficient number of users in small domains. Hence, although the accuracy of the *local* and *remote-average* approaches is higher, they cannot generate the same number of recommendations as the *standard* and *cross-domain* approaches. This will negatively affect the ability of the system to recommend all the interesting movies.[5] Hence, to properly evaluate the performance of the proposed approaches, the *coverage* of the recommendations (Herlocker et al. 2004), i.e., the percentage of the recommendations that are predictable, should be considered in addition to the traditional MAE measure.

The experimental results show that for any percentage of movies rated in the target genre, both the *cross-domain* and the *remote-average* mediation approaches significantly outperform the *standard* approach. The better accuracy of the *cross-domain* approach can be explained by arguing that its similarity computation is more accurate than the *standard* similarity computation. This is achieved by the fact that the *cross-domain* approach integrates users' domain-related similarities in a weighted manner, according to the inter-domain distances, whereas the *standard* approach disregards the inter-domain distances and assigns equal weights to all the available ratings. Hence, the similarity computation of the *cross-domain* approach is more accurate, as it performs a non trivial mediation of the user modelling data, whereas the standard approach performs a simple merge of the data. As a result, the MAE of the recommendations for the *standard* approach is lower.

The better accuracy of the *remote-average* recommendations can be explained by arguing that this mediation approach generates the recommendations independently by averaging the recommendations from several domains to which the movie belongs. In other words, the *remote-average* approach is a variant of nearest-neighbor multiple classifier system and this result is another example of the advantages of such a general approach (Kittler et al. 1998). As such, the similarity computation of the *remote-average* approach, which is using the ratings of related items only (i.e., items sharing one of the domains of the target item), yields more accurate similarity values than the *standard* similarity computation, which is using all the available ratings. Hence, the *remote-average* recommendations are more accurate than the *standard* ones.

As can be seen from the results, in some conditions *cross-domain* and *remote-average* mediations can be beneficial, as they can improve, respectively the coverage and the accuracy of the generated recommendations. However, deciding whether the mediation should be applied and choosing the most appropriate mediation approach is not a simple task. This decision depends on several factors, such as the overall goal and priorities of the recommendations (e.g., accuracy vs. coverage), sparsity of the avail-

[5] This deficiency of *local* and *remote-average* approaches was confirmed by the evaluation.

able user modeling data, availability of external domain-related data and inter-domain distances, and others.

4.3 Cross-context mediation

Cross-context UM mediation deals with importing and integrating past experiences of the same user for the same item in another context. Although cross-context mediation implies the contextual conditions to be included in the descriptions of the experiences, the state-of-the-art recommender systems mostly ignore it. As a result, currently there are no publicly available datasets that include descriptions of the contextual conditions of the experiences. Consequently, cross-context mediation has not been evaluated yet. Furthermore, several prior papers discussed context-aware recommender systems (Ricci 2006; Setten et al. 2004), context-aware collaborative filtering (Chen 2005; Lemire et al. 2005), and context-aware personalization in general (Keidl and Kemper 2004; Oku et al. 2007). However, due to the lack of publicly available context-aware datasets, they mostly motivated the need for context-awareness personalization, provided example personalization scenarios, or proposed various high-level conceptual frameworks, but did not evaluate them thoroughly.

Initial experimental evaluation of context-aware recommendations, which also included certain aspects of cross-context UM mediation, was reported in Adomavicius et al. (2005). The following section will analyze that paper in the context of the proposed user modelling data mediation framework and use it to demonstrate a simplified case of cross-context mediation. The experiments reported by Adomavicius et al. (2005), were conducted on a movies recommender system that included a specific component for collecting context-aware data, i.e., user experiences containing the contextual conditions. Three contextual features were recorded by the system: (1) *time*—specifies whether the movie was seen on a weekend, weekday, or it is unknown, (2) *place*—specifies whether the movie was seen in a theater, at home, or it is unknown, and (3) *companion*—specifies whether the movie was seen alone, with the family, with a boy/girlfriend, or with friends.

The above contextual features were represented by a three-dimensional contextual space. Each contextual feature could be assigned only one predefined value: *time* could be weekday, weekend, or unknown, *place* could be home, theater, or unknown and *companion* could be alone, with a boy/girlfriend, with family or with friends. Hence, every dimension was virtually split into intervals representing the possible values. Figure 10 illustrates the three-dimensional contextual space reflecting the above three features and their respective values. Every combination of the feature values produced a three-dimensional sub-space, referred to as *contextual segment*. For example, the above three features of *time*, *space*, and *companion*, and their respective values allow partitioning the space into $3 \times 3 \times 4 = 36$ contextual segments. To enhance flexibility of the system, contextual segments might be grouped into *aggregated segments*. For example, the highlighted segment in Fig. 10 shows the aggregated segment of being *alone* of the above three-dimensional contextual space.

All the recorded experiences were collaborative filtering ratings to movies given on a scale between 1 and 13. The ratings between 10 and 13 were considered *good*, while the ratings below 10 were considered *bad*. The ratings were virtually partitioned across

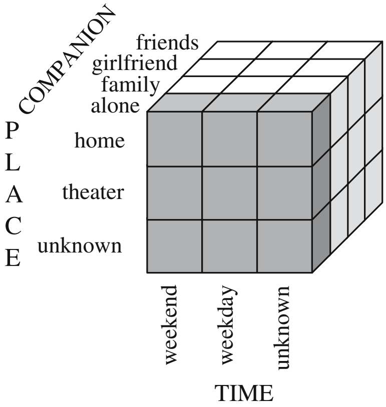

Fig. 10 Weekend contextual segment of the contextual space (adapted from Adomavicius et al. (2005))

the contextual segments, such that every segment stored only the experiences from the appropriate contextual conditions. As such, contextual conditions of the experiences stored in a certain segment were implicitly defined by the values of the contextual space dimensions. Hence, the experiences stored in every contextual segment were represented by standard two-dimensional collaborative filtering ratings matrices. In the rest of this section, the term *ratings matrix* refers to the experiences, i.e., user modeling data stored by the system.

Context-aware personalization was performed as a two-stage process. The first stage was referred to as *segment selection*. It focused on selecting the contextual segments (or aggregated segments), which store ratings matrices that should be used for the generation of recommendations. The second stage focused on generating collaborative filtering recommendations over the ratings matrices from the selected contextual segments. This process introduced a simplified variant of cross-context UM mediation, where only the available ratings matrices for specific, most appropriate contexts were mediated. Note that this type of mediation did not require integrating the imported ratings matrices, since user modeling data representation in the source and the target recommender systems was identical.

As part of the first segment selection stage, a fixed training set was constructed, and collaborative filtering recommendations were generated for the ratings in the training set. For every rating, two types of recommendations were generated: (1) using only the user modeling data stored in the ratings matrix of the contextual segment, to which the rating belonged, and (2) using the user modeling data stored in all the rating matrices, i.e., in all the segments of the contextual space. The accuracy of the generated recommendations in Adomavicius et al. (2005) was measured using the *F-measure* (Herlocker et al. 2004), considered as a representative example of decision-support accuracy measure, focusing on recommending only high-quality items. The *F-measure* was computed by:

$$F\text{-}measure = \frac{2 \cdot Precision \cdot Recall}{Precision + Recall}$$

Table 1 *F-measure* of the segment selection experiments

	Segment only	Entire dataset
Theater	0.608	0.479
Weekend	0.542	0.484
Theater-weekend	0.641	0.528

Table 2 *F-measure* of the recommendations

	Segment only	Entire dataset
Testing set	0.545	0.450

where, in terms of recommender systems, *Precision* was computed as the portion of the *good* (over 10) items ratings among the ratings that were recommended as *good* by the system, and, respectively, *Recall* was computed as the portion of the recommended *good* ratings among all the existing *good* ratings (Salton and McGill 1983).

Those contextual segments, where the accuracy of the first type of recommendations outperformed the accuracy of the second, were added to the set of selected segments. The selected segments were used at the second stage of generating collaborative filtering recommendations. In the experiments described in Adomavicuis et al. (2005), only 3 contextual segments were selected: (1) theater—segment of movies seen in a theater, (2) weekend—segment of movies seen at the weekend, and (3) theater-weekend—segment of movies seen at the weekend in a theater. The *F-measure* values of both types of the recommendations for the selected segments are shown in Table 1.

These results showed that in some contextual conditions (precisely, only in 3 contextual segments out of the 36 possible), the accuracy of the recommendations using the ratings matrix from the relevant segment was better than the accuracy of the recommendations using the ratings matrix of the entire dataset. In other words, the segment selection stage showed that, when requested to generate recommendations for a rating from one of these segments, the target recommender system should use only the user modeling data stored in the ratings matrix of the relevant segment, rather than all the available user modeling data.

As the above contextual segments were selected, the ability of the system to generate collaborative filtering recommendations for items from a new testing set was evaluated. Note that when requested to generate recommendations for ratings not from one of the selected 3 segments, the recommender system should use all the available user modeling data. Hence, the testing set contained only the ratings from these 3 segments. Also in this experiment the *F-measure* was used for evaluating the accuracy of the recommendations. The overall *F-measure* values for both types of recommendations are shown in Table 2.

This result showed that using the user modeling data, stored only in the ratings matrix of the relevant contextual segment, significantly improves the accuracy of the generated recommendations with respect to the recommendations using all the available user modeling data. Hence, taking into account the contextual information of the experiences allowed upgrading the accuracy of the generated recommendations, even despite the fact that the amount of the user modeling data, which was available to the recommendation generation stage, decreased.

However, it should be stressed that the accuracy improvement was not straightforward and it heavily depended on the available user modeling data. For example, the above segment selection stage showed that the accuracy improvement due to using the domain-specific ratings matrix was observed only in 3 segments out of 36 (in fact, in two segments and one aggregated segment). In these segments, the accuracy of the recommendations generated over the restricted user modeling data was improved, whereas in the rest of the contextual space the recommendations generated over all the available data were more accurate. It could be hypothesized that applying a more sophisticated learning mechanism at the segment selection stage, and further integrating user modeling data from several contextual segment may further improve the accuracy of the generated recommendations.

In summary, all the presented examples (cross-technique, cross-domain, and cross-context) demonstrated the potential of mediation of the user modeling data. They showed that in certain conditions, enriching the user modeling data available at the target recommender system may result in an improvement in the accuracy, and consequently, the quality of the generated recommendations. However, the improvement in the accuracy is not guaranteed and the decision regarding applying the mediation should be considered carefully.

5 Related works

The notion of general, i.e., application-independent UMs was initially proposed in Finin and Drager (1986). GUMS, a prototype of general user modeling system (Finin 1989), allowed the developers of personalization applications to define simple user stereotype hierarchies. For each stereotype, GUMS allowed facts that describe stereotype members and rules that facilitate reasoning about the stereotype members to be defined. At the runtime, GUMS received new facts about the user, verified their consistency with the currently available UMs, and answered queries concerning the assumptions about the user. Although GUMS has never been applied with real personalization applications, it determined the basic functionality of general user modeling systems: to constitute an independent component providing configurable user modeling services to the target personalization applications. Unlike GUMS, that suggested a general user modeling component that needs to be tailored and used by specific applications, this work proposes using any data collected by some recommender system and mediating them a form and representation required by the target system. Hence, no fixed and a priori defined structure and format are needed.

In Kobsa (2001), the authors surveyed the state-of-the-art research of general user modeling systems, and discussed the requirements that will facilitate their wide dissemination for both academic and commercial purposes:

- Generality—domain independence, compatibility with as many applications and domains as possible, and for as many user modeling tasks as possible.
- Expressiveness—ability to express as many types of facts and rules about the user as possible.
- Inferential capabilities—capability of performing various types of reasoning and resolving the conflicts when contradictory facts or rules are detected.

- Import of external data—ability to integrate the user modeling data collected by the system with the data collected by other systems, especially in distributed and ubiquitous environments.
- Privacy—support of privacy policies and conventions, national and international privacy legislations, and privacy-supporting tools and service providers.
- Quick adaptation—ability to quickly adapt services to new users, personalization functionalities, applications, and domains.
- Load balancing—keeping reasonable technical performance, i.e., robustness, high availability and low response times.

That paper also listed requirements from future general user modeling systems. They include ubiquitous modeling of mobile users, personalization functionalities in everyday smart appliances, and multi-purpose reuse of the general user modeling systems, not limited to the personalization purposes and expanding to other applications.

To facilitate reuse of user modeling data by multiple personalization applications, there is a need for semantically-enriched and ontology-based representation of the UMs. This issue was first discussed in Kay (1999), which motivated ontology-based reusable and understandable modeling of students. Reusability allowed separation of the representation of the UM from the personalization task and expansion of the notion of UMs limited to a particular application or application domain. The structure of the UMs was based on a set of predefined ontologies that facilitated access to a customized explanation of the meaning of the UM components in each domain.

The notion of ontology-based UMs was further developed in Razmerita et al. (2003), which presented a generic ontology-based user modeling architecture called OntobUM. OntobUM integrated three ontologies: user ontology characterizing the users, domain ontology defining the relationships between the personalization applications, and log ontology defining the semantics of user-application interaction. A similar approach for ontology-based representation of the UMs was presented in Heckmann et al. (2005). That paper introduced GUMO, a comprehensive set of General User Model Ontologies, which allowed uniform interpretation of distributed UMs in intelligent environments. GUMO was represented using OWL (Resource Description Framework), and was available for multiple personalization applications at the same time. Such commonly accepted ontology allowed the exchange of UM data between personalization applications to be simplified and the inherent problems of syntactical and structural differences between their UM representations to be overcome. The main constraint ontology-based approach is the inherent need for a commonly acceptable standard ontological representation. However, nowadays systems typically do not exploit a common user modeling ontology and agreed domain terminology. The approach proposed in this work is complementory to the ontology-based approach in the sense that it allows transformation of user modeling data between recommender systems even in a setting, where is no such common ontology. Moreover, it can be improved by incorporating a generic user modeling ontology, e.g., GUMO (Heckmann et al. 2005).

Several architectures for the exchange of user modeling data between multiple personalization applications in a distributed environment were discussed in previous studies. A centralized architecture for user modeling data exchange was presented and overviewed in Kay et al. (2003). That paper presented a setting, where a general UM stored by a central server was, in fact, a composition of partial UMs, stored by various personalization applications. Each application maintained its own inference mechanism that allowed it to update the general UM and to extract from it the needed information. When a UM was needed, the relevant information was extracted from the general UM and the inference mechanism adapted it to the needs of the target application. Thus, every application could generate its own view on the general UM. Differently from the approach proposed in Kay et al. (2003) implying existence of a centralized UM as a single source of user modeling data, this work proposes a semi-centralized and direct exchange of user modeling data between personalization applications. As this option does not imply any centralized management of the UMs, it is more dynamic and flexible.

A similar approach was discussed in Mehta et al. (2005), which proposed using a Unified User Context Model for improving the refinement of the collected UMs by exploiting parts of the same general ontology-based UM in several personalization applications. That paper discussed cross-application personalization approach, based on a notion of *context passport*. According to this approach, the target application extracted the required user modeling data from the context passport, provided the personalized services based on the extracted data, and updated the passport, if new user modeling data were available. Another outcome of that paper was the definition of the main stages of cross-application communication protocol:

- Negotiation—achieves an understanding on the type of user modeling data that is needed, i.e., agreeing on common ontology and vocabulary.
- Personalization—extraction of the user modeling data relevant to the activity and its transfer to the target application.
- Synchronization—replication and update of the stored user modeling data upon completion of personalization tasks.

Similarly to Kay et al. (2003), also Mehta et al. (2005) is based on a centralized storage of user modeling data.

A rather simplistic approach for resolving the problem of cross-application personalization through personal smart cards was suggested in Potonniee (2002). The smart cards stored and managed the UMs, partially resolving the privacy and availability issues, which are of great importance in distributed environments. Compared to a solution, where the profiles are stored in a central remote server, the use of smart cards made the profiles available in any context, enhanced privacy and security by allowing the users to control their own profiles, and avoided the communication delays. However, Potonniee (2002) discussed two-stage combination, which was applicable to collaborative filtering and content-based UMs, whereas the proposed approach provides a wider and generalized view of UM mediation framework, such that the above-mentioned combination of collaborative filtering and content-based UMs is only one possible option out of many others.

6 Conclusions and future research

This work presented a generic framework for UM mediation, which can be applied in recommender systems. The mediation is aimed at resolving the user modeling data sparseness problem, and, as a result, improving the quality of personalized context-aware recommendations provided to the users. Initially, this work discussed the user modeling data warehousing in various recommendation techniques and defined experience as a fundamental unit of the user modeling data. This definition encapsulated the representations of three primary dimensions of user, item, and context, and lead to a generic definition of UMs and user modeling data mediation. The user modeling mediation framework was then introduced, and several particular types of mediation were derived and discussed: cross-user mediation, cross-item mediation, cross-context mediation, and two types of cross-representation mediation.

Then, this work presented experimental results, for two mediation examples and an analysis of an existing paper demonstrating the potential of UM mediation for three cases presented: cross-technique (as a variant of cross-representation), cross-domain (as a generalized form of cross-item), and cross-context mediation. The results of cross-technique mediation demonstrated that mediating sparse collaborative filtering UMs to the appropriate content-based UMs may increase the accuracy of the generated recommendations. In cross-domain mediation of collaborative filtering user modeling data, the results demonstrated that a simple merge of domain-related UMs of the users or a weighted average of their domain-related similarity values do not improve the accuracy of the generated recommendations. Conversely, it is shown that mediating and averaging complete recommendations may improve it in a condition of very small UMs in the target domain. However, while accuracy is not easily improved, mediation can have a positive effect on the coverage of the recommendations, i.e., the capability of the system to provide reasonably good recommendations out of a set of potentially recommendable items. Also, the analysis of cross-context mediation demonstrated that in certain cases it may be beneficial to use only parts of the available user modeling data, storing past experiences in closely-related contextual conditions.

However, the conclusion regarding the benefit of the mediation of user modeling data is not straightforward. The evaluations demonstrated that the mediation may improve the quality of the generated recommendations and the performance of recommender systems only in certain conditions. These conditions are specific for each and every type of mediation and they include the type of mediated data, availability of user modeling data in the source and target systems, availability of external domain-related data, application domains of the involved systems and many other factors. Hence, the decision regarding applying the mediation should be taken only after a thorough analysis of these factors.

In summary, the main contribution of this work is in developing a generic mediation mechanism for integrating user modeling data in a distributed environment. The mediation mechanism facilitated interoperability of recommender systems and provision of better recommendations to the users. This contribution can be interpreted from two perspectives: (1) from the user modeling perspective, it established a novel approach to building accurate UMs by integrating user modeling data collected by a distributed set of recommender systems, and (2) from the recommender systems perspective, it provided a basis for a novel hybrid recommendation technique, where the recommendation process is based on multiple sources of user modeling data.

The discussed user modelling data mediation approaches mediation methods were quite simplistic and used intuitive reasoning mechanisms and shallow knowledge bases. In the future, it is planned to upgrade the accuracy of the mediation task by exploiting complex machine learning techniques for the reasoning task and applying various data mining techniques on the knowledge bases data. One data mining example may aim at developing domain ontology through analyzing domain-related textual data. The ontology will allow a better representation of user preferences. Another data mining example may aim at identifying similarities between users across several application domains and searching for behavioral patterns, which will allow generating predictions for users in certain domains basing on their behavior in other domains. This may further improve the accuracy of the provided recommendation and validate the proposed mediation of user modeling data.

Furthermore, several ongoing and future studies deal with practical implementation and evaluation of UM mediation in real-life systems from various application domains. For example, Berkovsky et al. (2007c) deals with UM mediation between Trip@dvice trip planning system (Ricci et al. 2002) and PIL personalized museum visitor's guide (Kuflik et al. 2007). The issues of cross-context mediation will be extensively studied within practical scenarios in the *SharedLife* project, focusing on multi-user shopping complemented by other everyday activities (Wahlster et al. 2006), and the *Passepartout* project focusing on search, browsing and viewing activities of individual and group users with a personalized digital TV guide (Aroyo et al. 2007).

Acknowledgements The authors gratefully acknowledge the support of the Caesarea Edmond Benjamin de Rothschild Foundation Institute for Interdisciplinary Applications of Computer Science (CRI) in Haifa, Israel, and of the Istituto Trentino di Cultura—the Center for Scientific and Technological Research (ITC-irst) in Trento, Italy.

References

Adomavicius, G., Sankaranarayanan, R., Sen, S., Tuzhilin, A.: Incorporating contextual information in recommender systems using a multidimensional approach. ACM Trans. Inform. Syst. **23**(1), 103–145 (2005)

Adomavicius, G., Tuzhilin, A.: Towards the next generation of recommender systems: a survey of the state-of-the-art and possible extensions. IEEE Trans. Knowledge Data Eng. **17**(6), 734–749 (2005a)

Adomavicius, G., Tuzhilin, A.: Personalization technologies: a process-oriented perspective. Commun. ACM **48**(10), 83–90 (2005b)

Aguzzoli, S., Avesani, P., Massa, P.: Collaborative case-based recommender system. In: European Conference on Case-Based Reasoning, pp. 460–474. Aberdeen, UK (2002)

Aroyo, L., Schut, H., Nack, F., Schiphorst, T., Kauw-A-Tjoe, M.: Personalized ambient media experience: move.me case study. In: International Conference on Intelligent User Interfaces, pp. 298–301. Honolulu, HI (2007)

Berkovsky, S.: Decentralized mediation of user models for a better personalization. In: International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, pp. 404–408. Dublin, Ireland (2006)

Berkovsky, S., Kuflik, T., Ricci, F.: Cross-technique mediation of user models. In: International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, pp. 21–30. Dublin, Ireland (2006a)

Berkovsky, S., Eytani, Y., Kuflik, T., Ricci, F.: Enhancing privacy and preserving accuracy of a distributed collaborative filtering. In: Recommender Systems Conference Minneapolis, MN, (2007c)

Berkovsky, S., Aroyo, L., Heckmann, D., Houben, G.J., Kröner, A., Kuflik, T., Ricci, F.: Providing context-aware personalization through cross-context reasoning of user modeling data. In: Workshop on Decentralized and Ubiquitous User Modeling, pp. 2–7. Corfu, Greece (2007a)

Berkovsky, S., Gorfinkel, A., Kuflik, T., Manevitz, L.: Case-based to content-based user model mediation. In: Workshop on Ubiquitous User Modeling, pp. 1–4. Riva del Garda, Italy (2006b)
Berkovsky, S., Kuflik, T., Ricci, F.: Distributed collaborative filtering with domain specialization. In: Recommender Systems Conference, Minneapolis, MN (2007b)
Brown, P.J., Bovey, J.D., Chen, X.: Context-aware applications: from the laboratory to the marketplace. IEEE Pers. Commun. **4**(5), 58–64 (1997)
Buriano, L., Marchetti, M., Carmagnola, F., Cena, F., Gena, C., Torre, I.: The role of ontologies in context-aware recommender systems. In: International Conference on Mobile Data Management, p. 80. Nara, Japan (2006)
Burke, R.: Knowledge-based recommender systems. In: Kent A. (ed.) Encyclopedia of Library and Information Systems, pp. 180–200. Dekker Encyclopedias (2000)
Burke, R.: Hybrid recommender systems: survey and experiments. User Model. User-Adapt. Interact. **12**(4), 331–370 (2002)
Chen, A.: Context-aware collaborative filtering system: predicting the user's preference in the ubiquitous computing environment. In: International Workshop on Location- and Context-Awareness, pp. 244–253. Oberpfaffenhofen, Germany (2005)
Claypool, M., Gokhale, A., Miranda, T., Murnikov, P., Netes, D., Sartin, M.: Combining content-based and collaborative filters in an online newspaper. In: Workshop on Recommender Systems, Berkeley, CA, http://www.cs.umbc.edu/~ian/sigir99-rec/ (1999)
Cranor, L.F., Reagle, J., Ackerman, M.S.: Beyond concern: understanding net users' attitudes about online privacy. Technical Report, AT&T Labs-Research (1999)
Dey, A.K., Abowd, G.D.: Towards a better understanding of context and context-awareness. In: International Symposium on Handheld and Ubiquitous Computing, pp. 304–307. Karlsruhe, Germany (1999)
Finin, T.W., Drager, D.: A general user modeling system. In: Canadian Conference on Artificial Intelligence, pp. 24–29. Montreal, Canada (1986)
Finin, T.W.: GUMS: a general user modeling shell. In: Kobsa, A., Wahlster, W. (eds.) User Models in Dialog Systems, pp. 411–430. Springer Publishers (1989)
Francisco-Revilla, L., Shipman, M.: Managing conflict in multi-model adaptive hypertext. In: Conference on Hypertext and Hypermedia, pp. 237–238. Santa Cruz, CA (2004)
Goker, A., Myrhaug, H.I.: User context and personalization. In: Workshop on Case Based Reasoning and Personalization, pp. 4–7. Aberdeen, UK (2002)
Goldberg, K., Roeder, T., Gupta, D., Perkins, C.: Eigentaste: a constant time collaborative filtering algorithm. Inform. Retr. **4**(2), 133–151 (2001)
Goldberg, D., Nichols, D., Oki, B., Terry, D.: Using collaborative filtering to weave an information tapestry. Commun. ACM **35**(12), 61–70 (1992)
Good, N., Schafer, J.B., Konstan, J.A., Borchers, A., Sarwar, B. Herlocker, J., Riedl, J.: Combining collaborative filtering with personal agents for better recommendations. In: National Conference on Artificial Intelligence, pp. 439–446. Orlando, FL (1999)
Hanani, U., Shapira, B., Shoval, P.: Information filtering: overview of issues, research and systems. User Model. User Adapt. Interact. **11**(3), 203–259 (2001)
Heckmann, D., Schwartz, T., Brandherm, B., Schmitz, M., Wilamowitz-Moellendorff, M.V.: GUMO—the general user model ontology. In: International Conference on User Modeling, pp. 428–432. Edinburgh, UK (2005)
Herlocker, J.L., Konstan, J.A., Borchers, A., Riedl, J.: An algorithmic framework for performing collaborative filtering. In: International Conference on Research and Development in Information Retrieval, pp. 230–237. Berkeley, CA (1999)
Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating collaborative filtering recommender systems. ACM Trans. Inform. Syst. **22**(1), 5–53 (2004) Internet Movie Database, http://www.imdb.com
Joachims, T., Freitag, D., Mitchell, T.: WebWatcher: a tour guide for the World Wide Web. In: International Joint Conference on Artificial Intelligence, pp. 770–777. Nagoya, Japan (1997)
Keidl, M., Kemper, A.: Towards context-aware adaptable web services. In: International World-Wide Web Conference, pp. 55–65. New York, NY (2004)
Kay, J.: Ontologies for reusable and scrutable student model. In: Workshop on Workshop on Ontologies for Intelligent Educational Systems, pp. 72–77. Le Mans, France (1999)
Kay, J., Kummerfeld, B., Lauder, P.: Managing private user models and shared personas. In: Workshop on User Modeling in Ubiquitous Computing, pp. 1–11. Pittsburgh, PA (2003)

Kittler, J., Hatef, M., Duin, R.P.W., Matas, J.: On combining classifiers. IEEE Trans. Pattern Anal. Machine Intell. **20**(3), 226–239 (1998)

Kobsa, A.: Generic User Modeling Systems. User Model. User-Adapt. Interact. **11**(1–2), 49–63 (2001)

Kobsa, A., Koenemann, J., Pohl, W.: Personalized hypermedia presentation techniques for improving online customer relationships. Knowledge Eng. Rev. **16**(2), 111–155 (2001)

Kohavi, R., John, G.H.: Wrappers for feature subset selection. Artif. Intell. J. **97**(1–2), 273–324 (1997)

Krulwich, B.: Lifestyle finder: intelligent user profiling using large-scale demographic data. Artif. Intell. Mag. **18**(2), 37–45 (1997)

Kuflik, T., Sheidin, J., Jbara, S., Goren-Bar, D., Soffer, P., Stock, O., Zancanaro M.: Supporting small groups in the museum by context-aware communication services. In: International Conference on Intelligent User Interfaces, pp. 305–308. Honolulu, HI (2007)

Lemire, D., Boley, H., McGrath, S., Ball, M.: Collaborative filtering and inference rules for context-aware learning object recommendation. J. Interact. Technol. Smart Educ. **2**(3), 179–188 (2005)

Linden, G., Smith, B., York, J.: Amazon.com recommendations: item-to-item collaborative filtering. IEEE Internet Comput. **7**(1), 76–80 (2003)

Maes, P.: Agents that reduce work and information overload. Commun. ACM **37**(7), 31–40 (1994)

Manouselis, N., Sampson, D.: A multi-criteria model to support automatic recommendation of e-learning quality approaches. In: World Conference on Educational Multimedia, Hypermedia and Telecommunications, pp. 518–526. Chesapeake, VA (2004)

McJones, P.: EachMovie Collaborative Filtering Data Set. HP Research (1997)

McNee, S.M., Lam, S., Konstan, J., Riedl, J.: Interfaces for eliciting new user preferences in recommender systems. In: International Conference on User Modeling, pp. 178–183. Pittsburgh, PA (2003)

Mehta, B., Niederée, C., Stewart, A., Degemmis, M., Lops, P., Semeraro, G.: Ontologically enriched user profiling for cross system personalization. In: International Conference on User Modeling, pp. 119–123. Edinburgh, UK (2005)

Montaner, M., Lopez, B., de la Rosa, J.L.: A taxonomy of recommender agents on the internet. AI Rev. **19**(4), 285–330 (2003)

Morita, M., Shinoda, Y.: Information filtering based on user behavior analysis and best match retrieval. In: International Conference on Research and Development in Information Retrieval, pp. 272–281. Dublin, Ireland (1994)

Mulvenna, M.D., Anand, S.S., Buchner, A.G.: Personalization on the net using web mining. Commun. ACM **43**(8), 123–125 (2000)

Oku, K., Nakajima, S., Miyazaki, J., Uemura, S.: Investigation for designing of context-aware recommendation system using SVM. In: International MultiConference of Engineers and Computer Scientists, pp. 970–975. Hong Kong, Hong Kong (2007)

Pascoe, J.: Adding generic contextual capabilities to wearable computers. In: International Symposium on Wearable Computers, pp. 92–99. Pittsburgh, PA (1998)

Pazzani, M.: A framework for collaborative, content-based and demographic filtering. Artif. Intell. Rev. **13**(5–6), 393–408 (1999)

Pazzani, M., Billsus, D.: Learning and revising user profiles: the identification of interesting web sites. Mach. Learn. **27**(3), 313–331 (1997)

Potonniee, O.: A decentralized privacy-enabling TV personalization framework. In: European Conference on Interactive Television, pp. 57–66. Brighton, UK (2002)

Rabanser, U., Ricci, F.: Recommender systems: do they have a viable business model in e-tourism? In: Conference of the International Federation for IT & Travel and Tourism, pp. 160–171. Lausanne, Switzerland (2005)

Razmerita, L., Angehrn, A., Maedche, A.: Ontology-based user modeling for knowledge management systems. In: International Conference on User Modeling, pp. 213–217. Pittsburgh, PA (2003)

Resource Description Framework, http://www.w3c.org/RDF/

Resnick, P., Iacovou, N., Sushak, M., Bergstrom, P., Riedl, J.: GroupLens: an open architecture for collaborative filtering of Netnews. In: International Conference on Computer-Supported Collaborative Work, pp. 175–186. Chapel Hill, NC (1994)

Resnick, P., Varian, H.R.: Recommender systems. Commun. ACM **40**(3), 56–58 (1997)

Ricci, F., Arslan, B., Mirzadeh, N., Venturini, A.: ITR: a case-based travel advisory system. In: European Conference on Case-Based Reasoning, pp. 613–627. Aberdeen, UK (2002)

Ricci, F.: Recommendations in context. In: International Conference on Electronic Commerce and Web Technologies, Krakow, Poland, keynote talk (2006)

Ricci, F., Cavada, D., Mirzadeh, N., Venturini, A.: Case-based travel recommendations. In: Fesemaier, D.R., Werthner, H., Wöber, K.W. (eds.) Destination Recommendation Systems: Behavioural Foundations and Applications, pp. 67–93. CABI publishers (2006)

Sakagami, H., Kamba, T.: Learning personal preferences on online newspaper articles for user behaviors. Comput. Networks ISDN Syst. **29**(8–13), 1447–1455 (1997)

Salton, G., McGill, M.: Introduction to Modern Information Retrieval. McGraw-Hill publishers (1983)

Sarwar, B., Karypis, G., Konstan, J., Riedl, J.: Item-based collaborative filtering algorithms. In: International World-Wide Web Conference, pp. 285–295. Hong Kong, Hong Kong (2001)

Schafer, J.B., Konstan, J., Riedl, J.: Electronic commerce recommender applications. J. Data Mining Knowledge Discov. **5**(1–2), 115–152 (2000)

Setten, M.V., Pokraev, S., Koolwaaij, J.: Context-aware recommendations in the mobile tourist application COMPASS. In: International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, pp. 235–244. Eindhoven, the Netherlands (2004)

Shardanand, U., Maes, P.: Social information filtering: algorithms for automating 'word of mouth'. In: International Conference on Human Factors in Computing Systems, pp. 210–217. Pittsburgh, PA, (1995)

Vozalis, M., Margaritis, K.G.: Enhancing collaborative filtering with demographic data: the case of item-based filtering. In: International Conference on Intelligent Systems Design and Applications, pp. 361–366. Budapest, Hungary (2004)

Wahlster, W., Kröner, A., Heckmann, D.: SharedLife: towards selective sharing of augmented personal memories. In: Sotck, O., Schaerf, M. (eds.) Reasoning, Action and Interaction in AI Theories and Systems, pp. 327–342. Springer Publishers (2006)

Wang, Y., Kobsa, A.: Impacts of privacy laws and regulations on personalized systems. In: Workshop on Privacy-Enhanced Personalization, pp. 44–46. Montreal, Canada (2006)

Ziegler, C.N., McNee, S.M., Konstan, J.A., Lausen, G.: Improving recommendation lists through topic diversification. In: International World-Wide Web Conference, pp. 22–32. Chiba, Japan (2005)

Authors' vitae

Shlomo Berkovsky is a Faculty of Department of Computer Science and Software Engineering at University of Melbourne. He graduated from the University of Haifa, where his research focused on mediation of user models in recommender systems. Currently, he is a research fellow at the Computer Science and Software Engineering Department at the University of Melbourne. His research project focuses on personalized content delivery for assisted navigation of information rich physical environments. His research interests include ubiquitous and distributed user modeling, recommender systems, context-aware personalization, and personalized text summarization.

Tsvi Kuflik is a Faculty of Department of Management Information Systems at University of Haifa. He received his BSc and MSc in Computer Science and PhD in Information Systems from the Ben-Gurion University of the Negev. Currently, he is a lecturer in the Management Information Systems Department at the University of Haifa. Before joining the University of Haifa he worked as a researcher at IRST research institute in Trento. At IRST, he worked on user modeling in active museums, work he continues at the University of Haifa in the framework of the collaboration projects of the University of Haifa and IRST, where he leads a research group focusing on supporting museum visitors. His research interests focus mainly on user modeling, ubiquitous user modeling and information filtering.

Francesco Ricci is a Faculty of Department of Computer Science at Free University of Bozen-Bolzano. He got his degree in Mathematics from University of Padova, and he is now associate professor at the Faculty of Computer Science, Free University of Bozen-Bolzano. His current research interests include recommender systems, intelligent interfaces, constraint satisfaction problems, machine learning and case-based reasoning. He is the author of many scientific publications appeared on international journals and conferences proceedings. He is in the editorial board of Information Technology and Tourism Journal and he is a member of ACM.