

Cross-domain Recommender Systems

Iván Cantador, Ignacio Fernández-Tobías, Shlomo Berkovsky, Paolo Cremonesi

Abstract The proliferation of e-commerce sites and online social media has allowed users to provide preference feedback and maintain profiles in multiple systems, reflecting a variety of their tastes and interests. Leveraging all the user preferences available in several systems or domains may be beneficial for generating more encompassing user models and better recommendations, e.g., through mitigating the cold-start and sparsity problems in a target domain, or enabling personalized cross-selling recommendations for items from multiple domains. Cross-domain recommender systems, thus, aim to generate or enhance recommendations in a target domain by exploiting knowledge from source domains. In this chapter, we formalize the cross-domain recommendation problem, unify the perspectives from which it has been addressed, analytically categorize, describe and compare prior work, and identify open issues for future research.

1 Introduction

Nowadays, the majority of recommender systems offer recommendations for items belonging to a single domain. For instance, Netflix recommends movies and TV programs, Barnes&Noble recommends books, and Last.fm recommends songs and music albums. These domain-specific systems have been successfully deployed by

Iván Cantador

Universidad Autónoma de Madrid, Madrid, Spain, e-mail: ivan.cantador@uam.es

Ignacio Fernández-Tobías

Universidad Autónoma de Madrid, Madrid, Spain, e-mail: ignacio.fernandezt@uam.es

Shlomo Berkovsky

CSIRO, Sydney, Australia, e-mail: shlomo.berkovsky@csiro.au

Paolo Cremonesi

Politecnico di Milano, Milan, Italy, e-mail: paolo.cremonesi@polimi.it

numerous websites, and the single-domain recommendation functionality is not perceived as a limitation, but rather pitched as a focus on a certain market.

Nonetheless, large e-commerce sites like Amazon and eBay often store user feedback for items from multiple domains, and in social media users often express their tastes and interests for a variety of topics. It may, therefore, be beneficial to leverage all the available user data provided in various systems and domains, in order to generate more encompassing user models and better recommendations. Instead of treating each domain (e.g., movies, books, and music) independently, knowledge acquired in a *source* domain could be transferred to and exploited in another *target* domain. The research challenge of transferring knowledge, and the business potential of delivering recommendations spanning across multiple domains, have triggered an increasing interest in *cross-domain recommendations*.

Consider two motivating use cases for cross-domain recommendations. The first refers to the well known cold-start problem, which hinders the recommendation generation due to the lack of sufficient information about users or items. In a cross-domain setting, a recommender may draw on information acquired from other domains to alleviate such problem, e.g., a user's favorite movie genres may be derived from her favorite book genres. The second refers to the generation of personalized cross-selling or bundle recommendations for items from multiple domains, e.g., a movie accompanied by a music album similar to the soundtrack of the movie. This recommendation may be informed by the user's movie tastes, but may not be extracted from rating correlations within a joined movie-music rating matrix.

These use cases are underpinned by an intuitive assumption that there are correspondences between user and item profiles in the source and target domains. This assumption has been validated in several marketing, behavioral, and data mining studies, which uncover strong dependencies between different domains [58, 66]. Cross-domain recommender systems leverage these dependencies through considering, for example, overlaps between the user or item sets, correlations between user preferences, and similarities of item attributes. Then, they apply a variety of techniques for enriching the knowledge of the target domain, and improving the quality of recommendations generated therein.

Cross-domain recommendation is a challenging and still largely under-explored topic. Although it has been studied from several angles, an agreed upon definition of the cross-domain recommendation problem has not emerged yet, and no work has analyzed and classified the existing cross-domain recommendation techniques. In this chapter we survey the state of the art in cross-domain recommender systems, categorize the methods for establishing and exploiting links between diverse domains, compare the outcomes of prior work, and outline future research directions.

The chapter is structured as follows. In Section 2 we formulate the cross-domain recommendation problem, describing its main tasks and goals. In Section 3 we present a general categorization of cross-domain recommendation techniques. In Sections 4 and 5 we review cross-domain recommendation approaches, distinguishing between knowledge aggregation and knowledge linkage/transfer approaches. In Section 6 we overview cross-domain recommendation evaluation. In Section 7 we

discuss practical considerations about cross-domain recommender systems. Finally, in Section 8 we discuss open research issues in cross-domain recommendation.

2 Formulation of the Cross-Domain Recommendation Problem

The cross-domain recommendation problem has been addressed from various perspectives in different research areas. It has been handled by means of user preference aggregation and mediation strategies for the cross-system personalization problem in user modeling [2, 8, 58], as a potential solution to mitigate the cold-start and sparsity problems in recommender systems [16, 59, 64], and as a practical application of knowledge transfer in machine learning [26, 40, 51].

Aiming to unify perspectives, we provide a generic formulation of the cross-domain recommendation problem, focusing on existing domain notions (Section 2.1) and cross-domain recommendation tasks (Section 2.2) and goals (Section 2.3), and discuss the possible scenarios of data overlap between domains (Section 2.4).

2.1 Definition of Domain

In the literature researchers have considered distinct notions of domain. For instance, some have treated items like *movies* and *books* as belonging to different domains, while others have considered items such as *action movies* and *comedy movies* as different domains. To the best of our knowledge, in the context of recommender systems research, there have been no attempts to define the concept of *domain*. Here we distinguish between several domain notions according to the attributes and types of recommended items. Specifically, we consider that *domain* may be defined at four levels (see illustration in Figure 1):

- *(Item) Attribute level*. Recommended items are of the same type, having the same attributes. Two items are considered as belonging to distinct domains if they differ in the value of certain attribute. For instance, two movies belong to distinct domains if they have different genres, like action and comedy movies. This definition of domain is rather borderline, and is mainly used as a way to increase the diversity of recommendations (e.g., we may wish to recommend some thriller movies even to users who only watch comedy movies).
- *(Item) Type level*. Recommended items are of similar types and share some attributes. Two items are considered as belonging to distinct domains if they have different attribute subsets. For instance, movies and TV shows belong to distinct domains, since although they have several attributes in common (title, genre), they still differ with respect to some others (e.g., the live attribute for TV shows).
- *Item level*. Recommended items are not of the same type, differing in most, if not all, of their attributes. For instance, movies and books belong to different domains, even though they have some attributes in common (title, release/publication year).

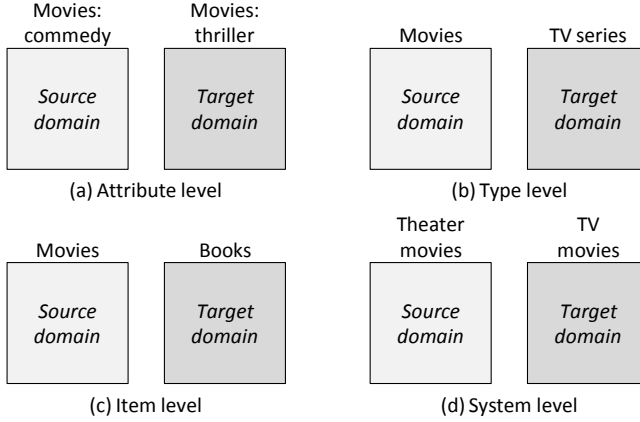


Fig. 1 Notions of domain according to attributes and types of recommended items. (a) *Attribute level*: same type of items (movies) with different values of certain attribute (genre). (b) *Type level*: similar types of items (movies and TV shows), sharing some of their attributes. (c) *Item level*: different types of items (books and movies). (d) *System level*: same type of items (movies) on different systems (theater and TV).

- *System level*. Recommended items belong to distinct systems, which are considered as different domains. For instance, movies rated in the MovieLens recommender, and movies watched in the Netflix video streaming service.

In Table 1 we summarize the considered notions of domains, addressed domains, and used datasets/systems in a significant number of prior works on cross-domain user modeling and recommendation. It can be seen that the majority of the papers considers domains at the item (about 55%) and system (24%) levels. The most frequently addressed domains are movies (75%), books (57%), music (39%) and TV (18%). In this context, we note that around 10% of the papers addresses various domains, by exploiting user preference data from multi-domain systems like Amazon and Facebook. Analyzing the pairs of domains frequently addressed, we observe that movies are often crossed with books (33%), music (19%), and TV (7%), whereas books are crossed with music (14%) and TV (10%).

The table also shows the utilized types of user preferences: ratings (61%), tags (29%), thumbs up (14%), transaction history (7%), and click-through data (4%). Although only a few papers use semantic concepts as user preferences, in some papers, social tags are transformed into concepts from WordNet or Wikipedia. Overall, about 14% of the papers use semantic-based user preferences.

2.2 Cross-Domain Recommendation Tasks

The research on cross-domain recommendation generally aims to exploit knowledge from a source domain \mathcal{D}_S to perform or improve recommendations in a target

Table 1 Summary of domain notions, domains, and user preference datasets/systems used in the cross-domain user modeling and recommendation literature.

Domain notion	Domains	User preferences - datasets/systems	References
<i>item attribute</i>	book categories	ratings - <i>BookCrossing</i>	Cao et al. 2010 [13]
	movie genres	ratings - <i>EachMovie</i>	Berkovsky et al. 2007 [7]
		ratings - <i>MovieLens</i>	Lee et al. 2001 [38] Cao et al. 2010 [13]
<i>item type</i>	books, movies, music	ratings - <i>Amazon</i>	Hu et al. 2013 [31] Loni et al. 2014 [44]
	books, games, music, movies & TV shows	ratings	Winoto & Tang 2008 [66]
<i>item</i>	books, movies	ratings - <i>BookCrossing</i> , <i>MovieLens/EachMovie</i>	Li et al. 2009 [40, 41] Gao et al. 2013 [26]
		ratings, tags - <i>LibraryThing</i> , <i>MovieLens</i>	Zhang et al. 2010 [67] Shi et al. 2011 [59] Enrich et al. 2013 [20]
		ratings, transactions	Azak 2010 [3]
		ratings - <i>Imhonet</i>	Sahebi & Brusilovsky 2013 [55]
		ratings - <i>Douban</i>	Zhao et al. 2013 [69]
	movies, music	thumbs up - <i>Facebook</i>	Shapira et al. 2013 [58]
	books, movies, music	tags - <i>MovieLens</i> , <i>Last.fm</i> , <i>LibraryThing</i>	Fernández-Tobías et al. 2013 [23]
		thumbs up - <i>Facebook</i>	Tiroshi & Kuflik 2012 [65] Cantador et al. 2013 [10] Tiroshi et al. 2013 [64]
	books, movies, music, TV shows		
	music, tourism	semantic concepts	Fernández-Tobías et al. 2011 [21] Kaminskas et al. 2013 [35]
	restaurants, tourism	ratings, transactions	Chung et al. 2007 [14]
	various domains	tags - <i>Delicious</i> , <i>Flickr</i>	Szomszor et al. 2008 [61, 62]
<i>system</i>	movies	ratings - <i>Netflix</i>	Cremonesi et al. 2011 [16] Zhao et al. 2013 [69]
		ratings - <i>Douban</i> , <i>Netflix</i>	Zhao et al. 2013 [69]
		ratings - <i>MovieLens</i> , <i>Moviepilot</i> , <i>Netflix</i>	Pan et al. 2012 [52] Pan et al. 2013 [53]
	music	tags - <i>Delicious</i> , <i>Last.fm</i>	Loizou 2009 [43]
		tags - <i>Blogger</i> , <i>Last.fm</i>	Stewart et al. 2009 [60]
	various domains	tags - <i>Delicious</i> , <i>Flickr</i> , <i>StumbleUpon</i> , <i>Twitter</i>	Abel et al. 2011 [1] Abel et al. 2013 [2]
		click-through data - <i>Yahoo!</i> services	Low et al. 2011 [45]




Fig. 2 Cross-domain recommendation tasks. Grey filled areas represent the target users and recommended items, and hatched areas represent the exploited data for generating recommendations.

domain \mathcal{D}_T . Analyzing the literature, we observe that the addressed tasks are diverse, and a consensual definition of the cross-domain recommendation problem has not been formulated yet. Hence, some researchers have proposed models aimed to provide jointly diverse recommendations of items belonging to multiple domains, whereas others have developed methods to alleviate cold-start and sparsity situations in a target domain by using information from source domains.

Aiming to provide a unified formulation of the cross-domain recommendation problem, we define the tasks we identify as providing recommendations across domains. Without loss of generality, we consider two domains \mathcal{D}_S and \mathcal{D}_T (the definitions are extensible to more source domains). Let \mathcal{U}_S and \mathcal{U}_T be their sets of users, and let \mathcal{I}_S and \mathcal{I}_T be their sets of items. The users of a domain are those who expressed preferences (e.g., ratings, reviews, tags, and consumption logs) for the domain items. The items do not necessarily have preferences from users of the domain, but may have content-based attributes that establish their membership to the domain.

Sorted in increasing order of complexity, we distinguish between the following three recommendation tasks (see Figure 2):

- *Multi-domain recommendation*: recommend items in both the source and target domains, i.e., recommend items in $\mathcal{I}_S \cup \mathcal{I}_T$ to users in \mathcal{U}_S (or, equivalently, in \mathcal{U}_T or $\mathcal{U}_S \cup \mathcal{U}_T$).
- *Linked-domain recommendation*: recommend items in the target domain by exploiting knowledge from the source and target domains, i.e., recommend items in \mathcal{I}_T to users in \mathcal{U}_S by exploiting knowledge about $\mathcal{U}_S \cup \mathcal{U}_T$ and/or $\mathcal{I}_S \cup \mathcal{I}_T$.
- *Cross-domain recommendation*: recommend items in the target domain by exploiting knowledge from the source domain, i.e., recommend items in \mathcal{I}_T to users in \mathcal{U}_S by exploiting knowledge about \mathcal{U}_S and/or \mathcal{I}_S .

Multi-domain approaches have mainly focused on the provision of cross-system recommendations, by jointly considering user preferences for items in various systems. To perform this type of recommendations, a significant overlap between user preferences in distinct domains is needed. This is becoming more and more feasible, since users maintain profiles in various social media, and there are interconnecting

mechanisms for both cross-system interoperability [12] and cross-system user identification [11]. In addition to social media, the benefits of multi-domain recommendations also come through in e-commerce sites, where personalized cross-selling [19, 36] can increase customer satisfaction, and consequently, their loyalty and the businesses profitability. For such purposes, in general, approaches aim to aggregate knowledge from the source and target domains.

Linked-domain approaches have been mainly explored to improve the recommendations in a target domain where there is a scarcity of user preferences, either at the user level (the cold-start problem) or at the community level (the data sparsity problem). To deal with these situations, a common solution is to enrich or enhance the available knowledge in the target domain with knowledge from the source domain. Hence, to perform this type of recommendations, some data relations or overlaps between domains are needed, and approaches aim to establish explicit or implicit knowledge-based links between the domains.

Finally, cross-domain approaches have been proposed to provide recommendations in a target domain where there is no information about the users. In this case, there is no assumption of data relations and/or overlaps between domains, and approaches aim to establish knowledge-based links between domains or to transfer knowledge from the source domain to the target domain.

For the sake of simplicity, we consider the three recommendation tasks together, as a single formulation of the cross-domain recommendation problem, although in Sections 4 and 5 we review specific approaches for each task.

2.3 Cross-Domain Recommendation Goals

From both the research and practical perspectives, it is important to match the recommendation algorithms to the task in hand. For this reason, we initially present a taxonomy of cross-domain recommendation goals. The taxonomy is described in a solution-agnostic way: each problem is defined based solely on its goals – without discussing how they are solved, which will be done in Section 3.

At the first level of the taxonomy, we consider the three recommendation tasks presented in Section 2.2, namely *multi-domain*, *linked-domain*, and *cross-domain* tasks, which are the columns of Table 2. At the second level, we distinguish between the specific goals addressed by cross-domain recommenders, which are the rows of Table 2. We identify the following goals:

- *Addressing the system cold-start problem (system bootstrapping)*. This is related to situations in which a recommender is unable to generate recommendations due to an initial lack of user preferences. One possible solution is to bootstrap the system with preferences from another source outside the target domain.
- *Addressing the new user problem*. When a user starts using a recommender, this has no knowledge of the user’s tastes and interests, and cannot produce personalized recommendations. This may be solved by exploiting the user’s preferences collected in a different source domain.

Table 2 Summary of cross-domain recommendation approaches based on goals and tasks.

<i>Goal</i>	<i>Multi-domain task</i>	<i>Linked-domain task</i>	<i>Cross-domain task</i>
<i>Cold start</i>			Shapira et al. 2013 [58]
<i>New user</i>		Winoto et al. 2008 [66] Cremonesi et al. 2011 [16] Low et al. 2011 [45] Hu et al. 2013 [31] Sahebi et al. 2013 [55]	Berkovsky et al. 2007 [6, 7] Berkovsky et al. 2008 [8] Nakatsuji et al. 2010 [47] Cremonesi et al. 2011 [16] Tiroshi et al. 2012 [65] Braunhofer et al. 2013 [9]
<i>New item</i>			Kaminskas et al. 2013 [35]
<i>Accuracy</i>	Cao et al. 2010 [13] Zhang et al. 2010 [67] Li et al. 2011 [42] Tang et al. 2011 [63] Zhang et al. 2013 [68]	Li et al. 2009 [40, 41] Moreno et al. 2012 [46] Shi et al. 2011 [59] Pan et al. 2012 [52] Gao et al. 2013 [26] Pan et al. 2013 [53] Zhao et al. 2013 [69]	Pan et al. 2008 [48] Stewart et al. 2009 [60] Pan et al. 2010 [51] Tiroshi et al. 2013 [64] Loni et al. 2014 [44]
<i>Diversity</i>		Winoto et al. 2008 [66]	
<i>User model</i>		Szomszor et al. 2008 [61] Abel et al. 2011 [1] Abel et al. 2013 [2] Fernández-Tobías 2013 [23] Goga et al. 2013 [28] Jain et al. 2013 [32]	

- *Addressing the new item problem (cross-selling of products).* When a new item is added to a catalog, it has no prior ratings, so it will not be recommended by a collaborative filtering system. This problem is particularly evident when cross-selling new products from different domains.
- *Improving accuracy (by reducing sparsity).* In many domains, the average number of ratings per user and item is low, which may negatively affect the quality of the recommendations. Data collected outside the target domain can increase the rating density, and thus may upgrade the recommendation quality.
- *Improving diversity.* Having similar, redundant items in a recommendation list may not contribute much to the user’s satisfaction (Chapter ??). The diversity of recommendations can be improved by considering multiple domains, as this may provide a better coverage of the range of user preferences.
- *Enhancing user models.* The main goal of cross-domain user modeling applications is to enhance user models. Achieving this goal may have personalization-oriented benefits such as (i) discovering new user preferences for the target domain [60, 62], (ii) enhancing similarities between users and items [1, 8], and (iii) measuring vulnerability in social networks [28, 32].

Table 2 shows the mapping between the above recommendation tasks and goals. Cross-domain tasks are mainly used to address the cold start problem boosting data density, while linked-domain tasks are used to improve accuracy and diversity.

2.4 Cross-Domain Recommendation Scenarios

As discussed by Fernández-Tobías et al. [22], in the context of a cross-domain recommendation task, domains can be explicitly or implicitly linked by means of content-based (CB) or collaborative filtering (CF) characteristics associated with users and/or items, such as ratings, social tags, semantic relations, and latent factors.

Let $\mathcal{X}^{\mathcal{U}} = \{x_1^{\mathcal{U}}, \dots, x_m^{\mathcal{U}}\}$ and $\mathcal{X}^{\mathcal{I}} = \{x_1^{\mathcal{I}}, \dots, x_n^{\mathcal{I}}\}$ be the sets of characteristics utilized to represent the users and items, respectively. Two domains \mathcal{D}_S and \mathcal{D}_T are linked if $\mathcal{X}_S^{\mathcal{U}} \cap \mathcal{X}_T^{\mathcal{U}} \neq \emptyset$ or $\mathcal{X}_S^{\mathcal{I}} \cap \mathcal{X}_T^{\mathcal{I}} \neq \emptyset$, i.e., if they share user or item characteristics. In a realistic setting, due to the heterogeneity of domain representations, one may need to set functions that map characteristics between domains, i.e., $f: \mathcal{X}_S^{\mathcal{U}} \rightarrow \mathcal{X}_T^{\mathcal{U}}$ and $g: \mathcal{X}_S^{\mathcal{I}} \rightarrow \mathcal{X}_T^{\mathcal{I}}$. For instance, to link movies and books, a mapping function could identify users registered in two systems, $f(u_i, \text{movie system}) = u_j, \text{book system}$, or could link related genres, $g(\text{comedy}_{\text{movie system}}) = \text{humor}_{\text{book system}}$.

Next, we describe representative examples of user and item characteristics, as well as their inter-domain relations and data overlap scenarios.

- *Content-based relations between domains.* In CB systems, a set of content or metadata features $\mathcal{F} = \{F_1, \dots, F_n\}$ – e.g., keywords, properties, and categories – describes both user preferences and item attributes, i.e., $\mathcal{X}^{\mathcal{U}} \subseteq \mathcal{F}, \mathcal{X}^{\mathcal{I}} \subseteq \mathcal{F}$. In general, a user profile is composed of a vector, where each component reflects the degree to which the user likes or is interested in a specific feature, and similarly, an item profile is composed of a vector whose components reflect the relevance of the features to the item. An overlap between domains \mathcal{D}_S and \mathcal{D}_T occurs when $\mathcal{X}_S^{\mathcal{U}} \cap \mathcal{X}_T^{\mathcal{U}} \neq \emptyset$ and $\mathcal{F}_S \cap \mathcal{F}_T \neq \emptyset$.
- *Collaborative filtering-based relations between domains.* In CF systems, user preferences are modeled as a matrix $R \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{I}|}$, in which an element $r_{u,i}$ is the rating assigned by user u to item i . Thus, $\mathcal{X}^{\mathcal{U}} = \mathcal{I}$ (\mathcal{I} being the rated items), and domains \mathcal{D}_S and \mathcal{D}_T overlap when $\mathcal{X}_S^{\mathcal{U}} \cap \mathcal{X}_T^{\mathcal{U}} \neq \emptyset$, i.e., $\mathcal{I}_S \cap \mathcal{I}_T \neq \emptyset$. An equivalent reasoning can be done for items, to derive that $\mathcal{X}^{\mathcal{I}} = \mathcal{U}$ (\mathcal{U} being the users with ratings), and that \mathcal{D}_S and \mathcal{D}_T overlap when $\mathcal{X}_S^{\mathcal{I}} \cap \mathcal{X}_T^{\mathcal{I}} \neq \emptyset$, i.e., $\mathcal{U}_S \cap \mathcal{U}_T \neq \emptyset$.

Moreover, as explained in subsequent sections, approaches have been proposed to represent users and/or items in lower dimension spaces, called *latent factors*, in which the above vector representations are valid. In these cases, if \mathbf{U} and \mathbf{I} denote the sets of user and item latent factors, respectively, then $\mathcal{X}^{\mathcal{U}} = \mathbf{U}$ and $\mathcal{X}^{\mathcal{I}} = \mathbf{I}$.

As shown in Figure 3, for the above types of relations, and generalizing the possible cross-domain CF cases identified by Cremonesi et al. [16], four scenarios of data overlap between two domains \mathcal{D}_S and \mathcal{D}_T can exist:

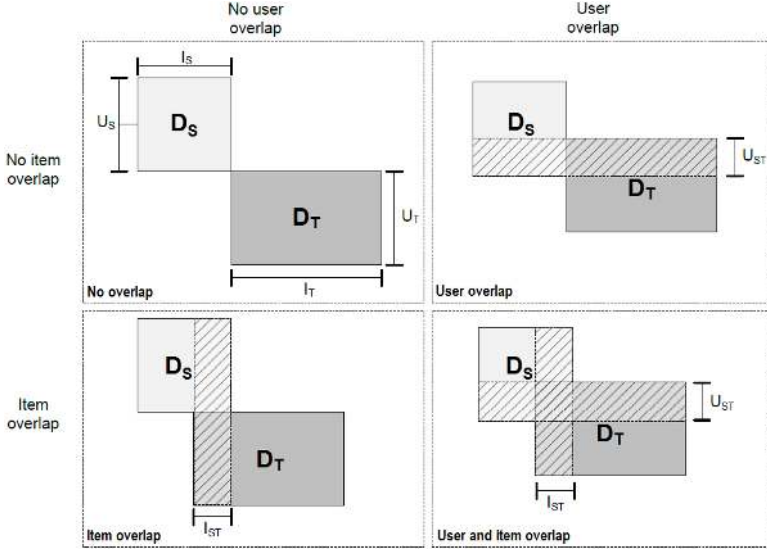


Fig. 3 Scenarios of data overlap between user and item sets in two domains \mathcal{D}_S and \mathcal{D}_T : *no overlap*, *user overlap*, *item overlap*, and *user and item overlap*.

- *No overlap*. There is no overlap between users and items in the domains, i.e., $\mathcal{U}_{ST} = \mathcal{U}_S \cap \mathcal{U}_T = \emptyset$ and $\mathcal{I}_{ST} = \mathcal{I}_S \cap \mathcal{I}_T = \emptyset$.
- *User overlap*. There are some common users who have preferences for items in both domains, i.e., $\mathcal{U}_{ST} \neq \emptyset$, but every item belongs to a single domain. This is the case, for instance, where some users rated movies and books.
- *Item overlap*. There are some common items that have been rated by users from both domains, i.e., $\mathcal{I}_{ST} \neq \emptyset$. This is the case, for instance, where two IPTV providers share a catalog of TV programs, which may be rated in each system.
- *User and item overlap*. There is overlap between both the users and items, i.e., $\mathcal{U}_{ST} \neq \emptyset$ and $\mathcal{I}_{ST} \neq \emptyset$.

3 Categorization of Cross-Domain Recommendation Techniques

As discussed in Section 2, cross-domain recommendation has been addressed from various perspectives in distinct research areas. This has entailed the development of a wide array of recommendation approaches, which in many cases are difficult to compare due to the user preferences they use, the cross-domain scenario they deal with, and the algorithms and data on which they are based. Moreover, published reviews of the research literature and categorizations of existing approaches [16, 22, 33, 39] have not reflected the entire complexity of the space. In this section, we categorize and propose a unifying schema for the existing cross-domain recommendation techniques.

Chung et al. presented in their seminal research [14] a framework that provides integrated recommendations for items that may be of different types, and may belong to different domains. The framework accounts for three levels of recommendation integration: *single item type recommendations*, which consist of items of the same type, *cross item type recommendations*, which consist of items of different types that belong to the same domain, and *cross domain recommendations*, which consist of items whose types belong to different domains. The authors stated that integrated recommendations can be generated by following at least three approaches:

- General filtering: instantiating a recommendation model for multiple item types that may belong to different domains.
- Community filtering: utilizing ratings shared among several communities or systems that may deal with different item types and domains.
- Market basket analysis: applying data mining to extrapolate hidden relations between items of different types/domains and to build a model for item filtering.

In [43], Loizou identified three main trends in cross-domain recommendation research. The first focuses on compiling unified user profiles appropriate for cross-domain recommendations [29]. This is considered as an integration of domain-specific user models into a single, unified multiple-domain user model, which is subsequently used to generate recommendations. The second involves profiling user preferences through monitoring their interactions in individual domains [34], which can be materialized through agents that learn single-domain user preferences and gather them from multiple domains to generate recommendations. The third deals with combining (or mediating) information from several single-domain recommender systems [6]. A number of strategies for mediating single-domain CF systems were considered: exchange of ratings, exchange of user neighborhoods, exchange of user similarities, and exchange of recommendations.

Based on these trends, Cremonesi et al. surveyed and categorized cross-domain CF systems [16]. They enhanced Loizou's categorization by considering a more specific grouping of approaches:

- Extracting association rules from rating behavior in a source domain, and using extracted rules to suggest items in a target domain, as proposed by Lee et al. [38].
- Learning inter-domain rating-based similarity and correlation matrices, as proposed by Cao et al. [13] and Zhang et al. [67].
- Combining estimations of rating probability distributions in source domains to generate recommendations in a target domain, as proposed by Zhuang et al. [70].
- Transferring knowledge between domains to address the rating sparsity problem in a target domain, as proposed by Li et al. [40, 41] and Pan et al. [50, 51].

For the last group, Li presented a survey of transfer learning techniques in cross-domain CF [39]. There, Li proposed an alternative categorization based on types of domain. Specifically, the author distinguished between (i) *system domains* that are associated with different recommenders, and represent a scenario where the data in a target recommender are very sparse, while the data in related recommenders are abundant; (ii) *data domains* that are associated with multiple sources of heterogeneous data, and represent a scenario where user data in source domains (e.g., binary

ratings) can be obtained easier than in a target domain (e.g., five-star ratings); and (iii) *temporal domains* that are associated with distinct data periods, and represent a scenario where temporal user preference dynamics can be captured. For these categories, Li considered three recommendation strategies differing in the knowledge transferred between domains:

- Rating pattern sharing, which aims to factorize single-domain rating matrices utilizing user/item groups, encode group-level rating patterns, and transfer knowledge between domains through the encoded patterns [40, 41, 42].
- Rating latent feature sharing, which aims to factorize single-domain rating matrices using latent features, share latent feature spaces across domains, and transfer knowledge between domains through the latent feature matrices [50, 51, 52, 53].
- Domain correlating, which aims to factorize single-domain rating matrices using latent features, explore correlations between latent features in single domains, and transfer knowledge between domains through such correlations [13, 59, 67].

Pan and Yang identified in a survey of transfer learning for machine learning applications [49] three main questions to be faced: (i) *what* to transfer – which knowledge should be transferred between domains; (ii) *how* to transfer – which learning algorithms should be exploited to transfer the discovered knowledge; and (iii) *when* to transfer – in which situations the knowledge transfer knowledge is beneficial. Focusing on the what and how questions, Pan et al. proposed in [50] and [51] a two-dimensional categorization of transfer learning-based approaches for cross-domain CF. The first dimension takes the type of transferred knowledge into account, e.g., latent rating features, encoded rating patterns, and rating-based correlations and covariances. The second dimension considers the algorithm, and distinguishes between adaptive and collective approaches, assuming, respectively, the existence of rating data only in the source domain, and in both the source and target domains.

In a more recent survey, Fernández-Tobías et al. went beyond CF recommendations, taking into account approaches that establish cross-domain relationships not necessarily based on ratings [22]. They identified three directions to address the cross-domain recommendation problem. The first is through the integration of single-domain user preferences into a unified cross-domain user model, which implies aggregating user profiles from multiple domains (“compile unified profiles” in [43]), and the mediation of user models across domains (“profile through monitoring” in [43]). The second direction aims to transfer knowledge from a source domain to a target domain, and includes approaches that exploit recommendations generated for a source domain in a target domain (“mediating information” in [43]), and approaches based on transfer learning, surveyed in [39]. The third direction is about establishing explicit relations between domains, which may be based either on content-based relations between items or on rating-based relations between users/items. The authors then proposed a two-dimensional categorization of cross-domain recommendation approaches: (i) according to the type of relations between domains: *content-based relations* (item attributes, tags, semantic properties, and feature correlations) vs. *rating-based relations* (rating patterns, rating latent factors,

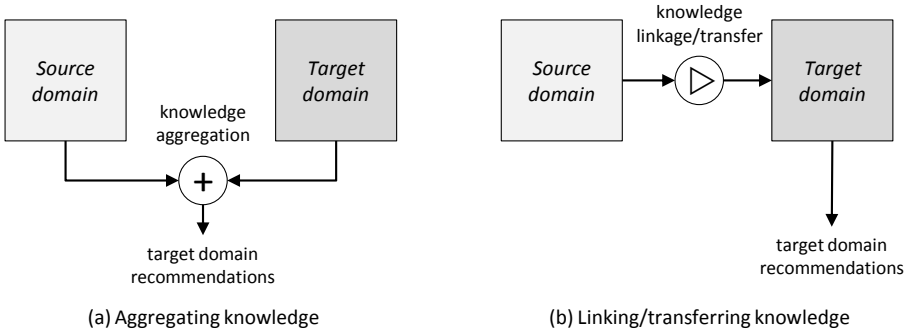


Fig. 4 Exploitation of knowledge in cross-domain recommendation.

and rating correlations); and (ii) according to the goal of the recommendation task: *adaptive models*, which exploit knowledge from a source domain to generate recommendations in a target domain, vs. *collective models*, which are built using data from several domains to improve recommendations in a target domain.

As can be seen from the previous discussion, the existing categorizations of cross-domain recommendation techniques are diverse. We aim to reconcile these categorizations in a way that captures and unifies their core ideas. For this, we focus on the exploitation of knowledge in cross-domain recommendation, which dictates the following two-level taxonomy:

- *Aggregating knowledge*. Knowledge from various source domains is aggregated to perform recommendations in a target domain (Figure 4a). Three use cases are considered, which will be analyzed in Section 4:
 - *Merging user preferences* – the aggregated knowledge consists of user preferences, e.g., ratings, tags, transaction logs, and click-through data.
 - *Mediating user modeling data* – the aggregated knowledge comes from user modeling data exploited by various recommender systems, e.g., user similarities and user neighborhoods.
 - *Combining recommendations* – the aggregated knowledge is composed of single-domain recommendations, e.g., rating estimations and rating probability distributions.
- *Linking and transferring knowledge*. Knowledge linkage or transfer between domains is established to support recommendations (Figure 4b). Three variants are considered, which will be analyzed in Section 5:
 - *Linking domains* – linking domains by a common knowledge, e.g., item attributes, association rules, semantic networks, and inter-domain correlations.
 - *Sharing latent features* – the source and target domains are related by means of implicit latent features.
 - *Transferring rating patterns* – explicit or implicit rating patterns from source domains are exploited in the target domain.

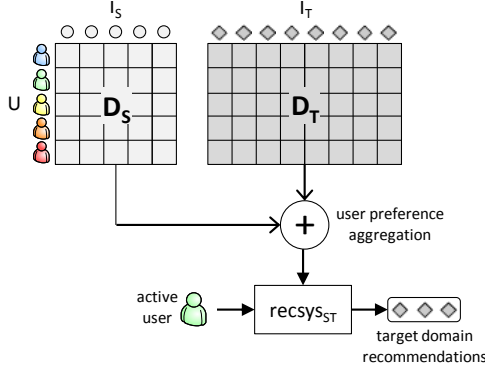


Fig. 5 *Merging user preferences.* Data sources from different domains are merged, and a traditional single-domain recommender system is used on the merged data.

4 Knowledge Aggregation for Cross-Domain Recommendations

In this section, we survey cross-domain recommendation approaches that aggregate knowledge from source domains to perform or improve recommendations in a target domain. The aggregated knowledge can be obtained at any stage of the recommendation process. In particular, it can be obtained from user preferences acquired at the user modeling stage (Section 4.1), from intermediate user modeling data utilized at the item relevance estimation stage (Section 4.2), or from item relevance estimations used at the recommendation generation stage (Section 4.3).

4.1 Merging Single-Domain User Preferences

Merging user preferences from different source domains is among the most widely used strategies for cross-system personalization, and the most direct way to address the cross-domain recommendation problem (see Figure 5).

Research has shown that richer profiles can be generated for users when multiple sources of personal preferences are combined, revealing tastes and interests not captured in isolated domains [2, 61]. It has been also shown that enriching sparse user preference data in a certain domain by adding user preference data from other domains, can significantly improve the generated recommendations under cold-start and sparsity conditions [55, 58]. These benefits, however, are accompanied by the need for having a significant amount of user preferences in multiple domains, and methods for accessing and merging the user profiles from different systems, which may have distinct types and/or representations of user preferences.

The most favorable scenario for aggregation-based methods implies that different systems share user preferences of the same type and representation. This scenario was addressed by Berkovsky et al. with a mediation strategy for cross-domain CF [6, 7]. The authors considered a domain-distributed setting where a global rating

matrix R is split, so that single-domain recommenders store local rating matrices R_d having the same structure. In this setting, a target domain recommender imports rating matrices R_d from the source domains, integrates the local and remote rating data into the unified rating matrix R , and applies CF to R . Note that this approach can be seen as a centralized CF with ratings split across multiple domains. Nonetheless, in this approach, smaller rating matrices are more efficiently maintained by local systems, and the data is shared with the target system only when needed.

Berkovsky et al. [6, 7] showed an improvement in the accuracy of target domain recommendations when aggregating ratings from several domains. This was also observed by Winoto and Tang [66]. The authors collected ratings for items in several domains and conducted a study that revealed that even when there exists significant overlap and correlation between domains, recommendation accuracy in the target domain is higher if only ratings in such domain are used. Despite these findings, Winoto and Tang stated that cross-domain recommendations may have alternative benefits, in particular, serendipity and diversity.

Apart from serendipity and diversity, other benefits of cross-domain recommendations have been identified. Sahebi and Brusilovsky [55] examined the impact of the size of user profiles in the source and target domains on the quality of CF, and showed that aggregating ratings from several domains allows increasing the accuracy of recommendations in the target domain under cold-start conditions. Similarly, Shapira et al. showed significant accuracy improvements by using aggregation-based methods when the available user preferences are sparse [58]. In this case, the authors used a dataset composed of unary Facebook *likes* as user preferences.

Beyond numeric ratings and unary/binary data, other types of user preferences have also been aggregated for cross-domain recommendations. In particular, several studies have focused on aggregating user profiles composed of social tags and semantic concepts. In this context, there is no need for user or item overlap between domains, since tags and concepts are used as a common representation to merge user preferences from multiple domains.

Szomszor et al. were among the first to correlate tag-based user profiles from multiple systems. In [62], they presented an architecture that transforms a set of raw tags into a set of filtered tags aligned between folksonomies in different domains. Crossing social-tag based profiles from the Delicious and Flickr folksonomies, the authors showed that filtered tags increase the overlap between domains, and allows discovering prominent user interests, locations, and events. In a follow-up work [61], Szomszor et al. extended their framework to map social tags to Wikipedia concepts, and build cross-domain user profiles composed of Wikipedia categories. An evaluation showed that new concepts of interest were learnt when expanding a user tag cloud with an external repository. Related to these works, Abel et al. [1] investigated the aggregation of a user's tag clouds from multiple systems. They evaluated a number of methods for semantic enrichment of tag overlap between domains, via tag similarities and via association rules deduced from the tagging data across systems. Aiming to analyze commonalities and differences among tag-based profiles, in [2] Abel et al. mapped tags to WordNet categories and DBpedia concepts. They used the mapped tags to build category-based user profiles, which revealed

significantly more information about the users than the profiles available in specific systems. Also in the context of tag-based user profile aggregation, Fernández-Tobías et al. [23] presented an approach that maps tags to emotional categories, under the assumption that emotions evoked by items in an entertainment domain can be represented through tags of folksonomies in which the items are annotated. Hence, emotions assigned to preferred items would be the bridge to merge user profiles across domains.

Regarding the use of semantic concepts as user preferences, Loizou [43] presented an approach that builds a graph where the nodes are associated with Wikipedia concepts describing items liked by the users, and the edges encode the semantic relationships between those concepts, obtained by integrating user ratings and Wikipedia hyperlinks. Using such a graph, a Markov chain model was used to produce recommendations by assessing the probability of traversing the graph towards a particular item, using the nodes in the user's profile as starting points. A related approach was studied by Fernández-Tobías et al. [21] and by Kaminskis et al. [35]. The authors presented a knowledge-based framework of semantic networks that link concepts from different domains. These networks are weighted graphs, in which nodes with no incoming edges represent concepts belonging to the source domain, and nodes with no outgoing nodes represent concepts belonging to the target domain. The framework provides an algorithm that propagates the node weights, in order to identify target concepts that are most related to the source concepts. Implemented on top of DBpedia, the framework was evaluated for recommending music suited to places of interest, which were related through concepts from several domains and contextual dimensions of location and time.

Instead of aggregating user preferences directly, several researches have focused on directed weighted graphs that link user preferences from multiple domains. In [47], Nakatsuji et al. presented an approach that builds domain-specific user graphs whose nodes are associated with users, and whose edges reflect rating-based user similarities. Domain graphs are connected via users who either rated items from several domains or shared social connections, to create a cross-domain user graph. Over this graph, a random walk algorithm retrieves items most liked by the users associated with the extracted nodes. Cremonesi et al. [16] built a graph whose nodes are associated with items and whose edges reflect rating-based item similarities. In this case, the inter-domain connections are the edges between pairs of items in different domains. The authors also proposed to enhance inter-domain edges by discovering new edges and strengthening existing ones, through strategies based on the transitive closure. Using the built multi-domain graph, several neighborhood- and latent factor-based CF techniques were evaluated. In [64], Tiroshi et al. collected a dataset containing user preferences in multiple domains extracted from social network profiles. The data was merged into a bipartite user-item graph, and various statistical and graph-based features of users and items were extracted from the graph. These features were exploited by a machine learning algorithm that addressed the recommendation problem as a binary classification problem.

Table 3 Summary of cross-domain user modeling and recommendation approaches based on merging single-domain user preferences.

Cross-domain approach	Inter-domain relationships	References
<i>Aggregating user ratings into a single multi-domain rating matrix</i>	Rating correlations	Berkovsky et al. 2007 [7] ^{UI} Sahebi & Brusilovsky 2013 [55] ^U Shapira et al. 2013 [58] ^U
	Rating correlations and relations between domain categories	Winoto & Tang 2008 [66] ^U
<i>Using a common representation for user preferences from multiple domains</i>	Social tag overlap	Szomszor et al. 2008 [62] ^N Szomszor et al. 2008 [61] ^N Abel et al. 2011 [1] ^N Abel et al. 2013 [2] ^N Fernández-Tobías et al. 2013 [23] ^N
	Semantic relationships between domain concepts	Loizou 2009 [43] ^N Fernández-Tobías et al. 2011 [21] ^N Kaminskas et al. 2013 [35] ^N
<i>Linking user preferences via a multi-domain graph</i>	Rating-based user/item similarities	Nakatsuji et al. 2010 [47] ^U Cremonesi et al. 2011 [16] ^U
	Patterns of user-item graph-based features	Tiroshi et al. 2013 [64] ^U
<i>Mapping user preferences to domain-independent features</i>	Socio-demographic and emotional features	González et al. 2005 [29] ^N
	Personality features	Cantador et al. 2013 [10] ^N
	User-item interaction features	Loni et al. 2014 [44] ^U

(N) no overlap, (U) user overlap, (I) item overlap, (UI) user and item overlap

The last type of cross-domain recommendation based on user preference aggregation is formed by the approaches that map user preferences from multiple domains to domain-independent features, and use the mapped feature-based profiles to build machine learning models that predict a user's preferences in the target domain. Although not conducting evaluations, González et al. [29] proposed an approach for unifying single-domain user models by interoperability and coordination of several agents. In addition to user tastes and interests, the unified model is composed of the user's socio-demographic and emotional features. Focusing on user personality features, Cantador et al. [10] studied the relations that exist between personality types and user preferences in multiple entertainment domains, namely movies, TV, music, and books. They analyzed a large number of Facebook user profiles composed of both Big Five personality trait scores [15] and explicit preferences for 16 genres in each of the above domains. As a result, the authors inferred similarities between personality-based user stereotypes in different domains. Finally, Loni et al. [44] presented an approach that encodes rating matrices from multiple domains as

real-valued feature vectors. With these vectors, an algorithm based on factorization machines [54] finds patterns between features from the source and target domains, and outputs preference estimations associated with the input vectors.

We summarize the discussed aggregation-based methods in Table 3. Aggregating ratings from several CF systems is the simplest method, but requires access to user profiles, and a significant rating overlap between domains, which may not be achievable in real situations. Thus, most aggregation-based methods transform user preferences from multiple domains into a common representation, independent of the domains of interest, and usable for establishing inter-domain data relations and overlaps. For this purpose, social tags and semantic concepts serve as the main types of user preferences. More recent methods focus on aggregating several sources of user preferences from multiple domains into a single graph. Due to the increasing use of social media, we envision that novel cross-domain recommendation approaches that both unify user preferences and aggregate them into multi-domain graphs will be developed.

4.2 Mediating Single-Domain User Modeling Data

Not only immediate user preferences, but also other recommendation-related information about users, items, and domains may be aggregated or mediated (see Figure 6). An early approach for cross-domain recommendation through mediation was proposed by Berkovsky et al. [8]. The central idea behind user model mediation is that importing any user modeling data from source recommenders may benefit a target recommender [4] – the mediation can enrich the user models of the target recommender, and yield more accurate recommendations. What data can be mediated between the source and the target recommenders? The most simple scenario covered in Section 4.1 includes importing the user models, whereas more complex scenarios include mediating specific recommendation data.

For example, in a CF system, cross-domain mediation may import the list of nearest neighbors. This is underpinned by two assumptions: (i) there is overlap of users between domains, and (ii) user similarity spans across domains, i.e., if two users are similar in a source domain, they are similar also in the target domain. This idea was leveraged in the heuristic variant of cross-domain mediation developed by Berkovsky et al. [7]. There, it was shown that importing nearest neighbors, and computing their similarity with the target domain data only, can produce more accurate recommendations than single-domain recommendations. A similar idea was formulated by Shapira et al. [58] as the k nearest neighbors (k -NN) source aggregation. They used multi-domain Facebook data to produce the set of candidate nearest neighbors, and compute their local similarity degree in the source domain. This allowed overcoming the new user problem and the lack of ratings in the target domain. Another attempt to use multi-domain Facebook data was done by Tiroshi and Kuflik [65]. They applied random walks to identify source domain-specific neighbor sets, which were used to generate recommendations in the target domain.




Fig. 6 Mediating user modeling data. A model is learnt in the source domain (e.g., the neighborhood of a user) and used in the target domain.

Aggregating the lists of nearest neighbors relies on their data in the target domain only, which may be too sparse and result in noisy recommendations. Thus, one could consider importing and aggregating also the degree of their similarity in the source domain. This approach was referred to in [7] as cross-domain mediation. A content-based and a statistical variant of domain distance metrics were evaluated in [5], producing comparable results and outperforming single-domain recommendations. The weighted k-NN aggregation was further enhanced by Shapira et al. [58]. The authors compared several weighting schemes, the performance of which was consistent across several metrics and recommendation tasks. The above scenarios of cross-domain mediation assume an overlap in the sets of users. An analogous scenario refers to a setting where items overlap between the source and target domains, which opens the opportunity for further mediation. One of them, involving only the music domain, but two systems (for tagging and for blogging) was studied by Stewart et al. [60]. The authors leveraged the tags assigned by similar users on Last.fm in order to recommend tags on Blogger.

Moving from CF to latent factor-based methods, we highlight two works compatible with the user modeling data mediation pattern. Low et al. [45] developed a hierarchical probabilistic model that combines user information across multiple domains, and facilitates personalization in domains with no prior user interactions. The model is underpinned by a global user profile based on a latent vector, and a set of domain-specific latent factors that eliminate the need for common items or features. Pan et al. [52] dealt with transferring uncertain ratings, i.e., expected rating range or distribution derived from behavioral logs, using latent features of both users and items. The uncertain ratings were transferred from the source to the target domain, and leveraged there as constraints for the matrix factorization model.

We summarize the mediation-based approaches in Table 4. As can be seen, they all imply either user- or item-overlap between the source and target domains. These are necessary for identifying high-level user preferences spanning across domains. This often requires sharing of user data between several systems, which is avoided

Table 4 Summary of cross-domain recommendation approaches based on mediating single-domain user modeling data.

Cross-domain approach	Inter-domain relationships	References
<i>Aggregating neiighbourhoods to generate recommendations</i>	Rating-based user similarities	Berkovsky et al. 2007 [7] ^U Tiroshi & Kuflik 2012 [65] ^{UI} Shapira et al. 2013 [58] ^U
<i>Aggregating user-to-user similarities to generate recommendations</i>	Content- and rating-based user similarities	Berkovsky et al. 2007 [7] ^U Shapira et al. 2013 [58] ^U
<i>Exploiting user neighborhoods to enhance target user models</i>	User overlap	Stewart et al. 2009 [60] ^I
<i>Combining probabilistic user models</i>	Latent features of domains and global user preferences	Low et al. 2011 [45] ^U
<i>Combining heterogeneous user preferences</i>	Domain-dependent constraints on matrix factorization	Pan et al. 2012 [52] ^{UI}

(N) no overlap, (U) user overlap, (I) item overlap, (UI) user and item overlap.

due to commercial competition and conflicts with privacy regulations. However, it is usual for a user to utilize multiple systems (or, in a more common use-case, to have accounts on multiple social networks), and thus cross-domain recommendations through mediation is a feasible scenario. Most of the surveyed approaches apply simple mediation methods, whereas the last two are based on latent representations, and apply probabilistic or transfer learning models. None of these works counts on explicit domain distance or similarity, which will be elaborated in Section 5.1). Hence, we conjecture that more future work will address cross-domain recommendation by mediating richer user modeling data.

4.3 Combining Single-Domain Recommendations

Overlap of both users and items allows aggregating ready-made single-domain recommendations (see Figure 7). Contrarily to the mediation-based cross-domain recommendation scenarios, the predicted recommendations from the source domain may inform on their own to the target domain recommender. Hence, the central question in combining single-domain recommendation refers to the weights assigned to recommendations coming from the source domains, which reflect their importance for the target domain. These weights may be computed through various factors, such as the reliability of each recommender, distance between the domains, and so forth.

The idea of combining single-domain recommendations was referred to in [6] and [7] as remote-average mediation. There, movie ratings were partitioned into do-




Fig. 7 Combining single-domain recommendations. Recommendations are generated independently for each domain and later merged for the final recommendation.

mains according to the genres of the movies. Since movies combine elements from multiple genres, and users watch movies from various genres, the user- and item-overlap are both present. This allows computing stand-alone recommendations in the source domains, and aggregating them for the target domain. Weighted aggregation of single-domain recommendations also was studied by Givon and Lavrenko [27]. The authors focused on the book recommendation task, accomplished using two different methods. Standard CF recommendations were complemented by relevance model-based recommendations, relying on the similarity of a book and the user's model, both consisting of book contents and tags assigned to the book. The two were combined in a weighted manner, such that the relative importance of the CF recommendations increased with the number of ratings available.

A relevant approach for cross-domain consensus regularization, although applied to classification problems and not to recommender systems, was proposed by Zhuang et al. [70]. The central contribution of that work is a framework for learning from multiple source domains, and reconciling discrepancies between the classifiers using the local data of the target domain. One of the advantages of the framework is that it does not require overlaps in either the user or item sets.

The overviewed approaches that combine single-domain recommendations are summarized in Table 5. Clearly, the key point for this group of cross-domain recommenders refers to the way the stand-alone source domain recommendations are combined. This is touched upon in [70], but also addressed in numerous researches outside the recommender systems space. It should be highlighted that the single-domain recommenders can use various techniques, and the combination of their outputs is independent of other components, e.g., user modeling, contextualization, and presentation, which makes this cross-domain aggregation variant attractive for practical recommenders.

Table 5 Summary of cross-domain recommendation approaches based on combining recommendations from single-domain user preferences.

Cross-domain approach	Inter-domain relationships	References
<i>Aggregating user rating predictions</i>	Rating-based user similarities	Berkovsky et al. 2007 [7] ^{UI} Givon & Lavrenko 2009 [27] ^{UI}
<i>Combining estimations of rating distribution</i>	Rating distribution similarities	Zhuang et al. 2010 [70] ^N

(*N*) no overlap, (*U*) user overlap, (*I*) item overlap, (*UI*) user and item overlap.

5 Knowledge Linkage and Transfer for Cross-Domain Recommendation

In this section, we survey cross-domain recommendation approaches that link or transfer knowledge between domains, enhancing the information available in the target domain for the generation of recommendations. The knowledge linkage and transfer can be done explicitly – e.g., via common item attributes, semantic networks, association rules, and inter-domain user preference similarities (Section 5.1) –, implicitly by means of latent features shared by domains (Section 5.2), or by means of rating patterns transferred between domains (Section 5.3).

5.1 Linking Domains

A natural approach to address the heterogeneity of several domains is to identify correspondences between their characteristics. For instance, we may link a particular movie and a book because both belong to genres that can be semantically mapped, e.g., comedy movies and humorous books. In general, such inter-domain correspondences may be established directly using some kind of common knowledge between domains, e.g., item attributes, semantic networks, association rules, and inter-domain preference-based similarities or correlations (see Figure 8).

These links are valuable sources of information for reasoning across domains. A recommender system could identify potentially relevant items in the target domain by selecting those that are related to others in the source domains, and for which the user has expressed a preference in the past. Besides, inter-domain similarities and correlations can be exploited to adapt or combine knowledge transferred from different domains. One of the earliest approaches for linking domains was explored by Chung et al. [14]. Aiming to support the decision making process in recommendation, they proposed a framework for designing personalized filtering strategies. In the framework, relevant items in the target domain are selected according to the attributes they have in common with items in the source domain the user is inter-




Fig. 8 *Linking domains.* An external knowledge source is used to link items from different domains. User preferences in the source domain may be used to adapt the item linkage.

ested in. That is, the inter-domain links are established through the overlap of item attributes, and no user or item overlap between the domains is required.

Conversely to the use case of [14], in a realistic setting, items are highly heterogeneous, and often no common attributes between domains can be found. To address this situation, we may establish more complex, likely indirect relations between items in different domains. Hence, when suitable knowledge repositories are available, concepts from several domains can be connected by the means of semantic properties, forming semantic networks that explicitly link the domains of interest. Along these lines, Loizou [43] proposed to use Wikipedia as a universal vocabulary to express and relate user preferences across multiple domains. The author presented an approach that builds a graph, the nodes of which represent concepts (Wikipedia pages) describing items liked by the users, and edges encode the semantic relationships between those concepts, obtained by integrating user ratings and Wikipedia hyperlinks. Using such a graph, a Markov chain model produces recommendations by assessing the probability of traversing the graph from the nodes in the user's profile as a starting point toward the recommendable items.

A major difficulty of the above approaches is the well known *knowledge acquisition* problem, that is, building the above mentioned knowledge repositories. To address this problem, information has to be extracted and stored in a formal and structured representation that can be exploited by a recommender. Fernández-Tobías et al. [21] and Kaminskis et al. [35] envisioned Linked Data as a solution to the problem. Specifically, they proposed a framework for extracting a multi-domain semantic network from the DBpedia ontology, which links items and concepts in the source and target domains. Over the extracted network, a constrained spreading algorithm computes semantic similarities to rank and filter items in the target domain.

Inter-domain association rules have also been explored as an alternative to relate various types of items. In this direction, Azak [3] presented a framework for cross-domain recommendation in which knowledge-based rules defined by domain experts facilitate mapping between attributes in distinct domains, e.g., “people who

like romance drama movies also like dramatic poetry books.” These rules are then used to enhance CB and CF recommendations, adjusting the predicted ratings whenever rule conditions hold. In [10], Cantador et al. related user personality types with domain-dependent preferences by means of automatically generated association rules. The authors also extracted personality stereotypes for sets of domain genres. Based on these stereotypes, inter-domain similarities were computed between genres, which may be used to support knowledge transfer between domains.

Instead of linking domains by mapping attributes, an alternative way to transfer knowledge is to compute similarities or correlations between domains based on user preference or item content analysis. In an early work, Berkovsky et al. [5] explored this idea aiming to identify related domains, from which user data would be imported and utilized to enrich the user model in the target domain. The proposed approach makes use of web directories to identify websites that characterize the domains of interest. Then, the approach establishes domain similarities by computing the cosine similarity between the TF-IDF term vectors of the domains’ websites. We note that this method requires no overlap of users or items, but rather an external source of representative documents classified to several domains.

Another way of exploiting inter-domain similarities for cross-domain recommendation consists of integrating them into the matrix factorization method [56]. Specifically, such similarities are imposed as constraints over user or item latent factors when jointly factorizing rating matrices. For instance, Cao et al. [13] proposed an approach in which inter-domain similarities are implicitly learnt from data, as model parameters in a non-parametric Bayesian framework. Since user feedback is used to estimate the similarities, user overlap between the domains is required. Addressing the sparsity problem, Zhang et al. [67] adapted the probabilistic matrix factorization method to include a probability distribution of user latent factors that encodes inter-domain correlations. One strength of this approach is that user latent factors shared across domains are not needed, allowing more flexibility in capturing the heterogeneity of domains. Instead of automatically learning implicit correlations in the data, Shi et al. [59] argued that explicit common information is more effective, and relied on shared social tags to compute cross-domain user-to-user and item-to-item similarities. Similarly to previous approaches, rating matrices from the source and target domains are jointly factorized; but in this case user and item latent factors from each domain are restricted, so that their product is consistent with the tag-based similarities.

We have reviewed approaches that establish links and compute similarities between domains, which are summarized in Table 6. We observe that the majority of the proposed methods do not require inter-domain user or item overlap. Instead, linking approaches exploit content information to establish the inter-domain relationships. Likewise, in [5] and [59], similarities are computed based on common text and social tags. For these approaches, it is also worth noticing that no one clearly outperforms the others, since most of them are designed for particular cross-domain scenarios and, to the best of our knowledge, have not been compared empirically.

Table 6 Summary of cross-domain user modeling and recommendation approaches based on linking domains.

Cross-domain approach	Inter-domain relationships	References
<i>Relating and filtering items via common attributes</i>	Item attribute overlap	Chung et al. 2007 [14] ^N
<i>Building semantic network linking domain concepts</i>	Semantic relationships between domain concepts	Loizou 2009 [43] ^N Fernández-Tobías et al. 2011 [21] ^N Kaminskas et al. 2013 [35] ^N
<i>Relating item types via knowledge-based rules</i>	Inter-domain knowledge-based rules	Azak et al. 2010 [3] ^N Cantador et al. 2013 [10] ^N
<i>Computing inter-domain similarities</i>	Text overlap	Berkovsky et al. 2006 [5] ^N
<i>Constraining matrix factorization with inter-domain similarities</i>	Rating overlap	Cao et al. 2010 [13] ^U Zhang et al. 2010 [67] ^N
	Social tag overlap	Shi et al. 2011 [59] ^N

(N) no overlap, (U) user overlap, (I) item overlap, (UI) user and item overlap.

5.2 Sharing Latent Features by Domains

Latent factor models are among the most popular CF techniques [37]. In these models, user preferences and item attributes, which are typically very sparse, are characterized through a reduced set of latent factors discovered from data, to obtain a denser representation. The assumption is that using the new representation, latent user preferences and item attributes are better captured and matched.

Related to the *what to transfer* aspect of transfer learning [49], latent factors shared between domains can be exploited to support cross-domain recommendations (see Figure 9). Also, as pointed in Section 3, two types of approaches have been studied to address the *how to transfer* aspect; namely, *adaptive* and *collective* models. In the former, latent factors are learnt in the source domain, and are integrated into a recommendation model in the target domain, while in the latter, latent factors are learnt simultaneously optimizing an objective function that involves both domains.

In [51], Pan et al. addressed the sparsity problem in the target domain following the adaptive approach, proposing to exploit user and item information from auxiliary domains where user feedback may be represented differently. In particular, they studied the case in which users express binary like/dislike preferences in the source domain, and utilize 1-5 ratings in the target domain. Their approach performs singular value decomposition (SVD) in each auxiliary domain, in order to separately compute user and item latent factors, which are then shared with the target domain. Specifically, transferred factors are integrated into the factorization of the rating

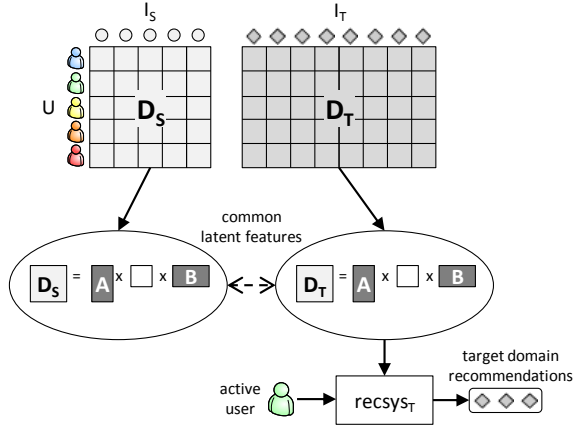


Fig. 9 *Sharing latent features.* Latent features models are learnt simultaneously on both the source and target domains, constraining user and/or item features to be the same across the domains.

matrix in the target domain and added as regularization terms so that specific characteristics of the target domain can be captured.

Latent factors can also be shared in a collective way, as studied by Pan et al. [50]. In this case, instead of learning latent features from the source domains and transferring them to the target domain, the authors proposed to learn the latent features simultaneously in all the domains. Both user and item factors are assumed to generate the observed ratings in every domain, and, thus, their corresponding random variables are shared between the probabilistic factorization models of each rating matrix. Moreover, the factorization method is further extended by incorporating another set of factors that capture domain-dependent information, resulting in a tri-factorization scheme. A limitation of the proposed approach is that the users and items from the source and target domains have to be identical.

Instead of focusing on sharing latent factors, Enrich et al. [20], and Fernández-Tobías and Cantador [24] studied the influence of social tags on rating prediction, as a knowledge transfer approach for cross-domain recommendations. The authors presented a number of models based on the SVD++ algorithm [37] to incorporate the effect of tag assignments into rating estimation. The underlying hypothesis is that information about item annotation in a source domain can be exploited to improve rating prediction in a target domain, as long as a set of common tags between the domains exists. In the proposed models, tag factors are added to the latent item vectors, and are combined with user latent features to compute rating estimations. The difference between these models is in the set of tags considered for rating prediction. In all the models knowledge transfer is performed through the shared tag factors in a collective way, since these are computed jointly for the source and the target domains.

In [31], Hu et al. presented a more complex approach that takes domain factors into account. There, the authors argue that user-item dyadic data cannot fully cap-

Table 7 Summary of cross-domain recommendation approaches based on latent features shared by domains.

Cross-domain approach	Inter-domain relationships	References
<i>Using user and item latent features of source domains to regularize latent features in a target domain</i>	Shared latent user preferences and latent item attributes	Pan et al. 2010 [51] ^{UI}
<i>Using the same latent factors to jointly factorize the rating matrices in the source and target domains</i>	User and item overlap	Pan et al. 2011 [50] ^{UI}
<i>Extending matrix factorization with a vector of latent factors associated to social tags</i>	Social tag overlap	Enrich et al. 2013 [20] ^N Fernández-Tobías & Cantador 2014 [24] ^N
<i>Sharing latent features via a user-item-domain tensor factorization</i>	Rating overlap	Hu et al. 2013 [31] ^U

(N) no overlap, (U) user overlap, (I) item overlap, (UI) user and item overlap.

ture the heterogeneity of items, and that modeling domain-specific information is essential to make accurate predictions in a setting, where users typically express their preferences in a single domain. They referred to this problem as the *unacquainted world*, and proposed a tensor factorization algorithm to exploit the triadic user-item-domain data. In that method, rating matrices from several domains are simultaneously decomposed into shared user, item, and domain latent factors, and genetic algorithm automatically estimates optimal weights of the domains.

Table 7 summarizes the described approaches sharing latent factors across domains. In contrast to the methods presented in Section 5.1, these approaches require inter-domain user or item overlap to extract shared latent factors, unless shared content information is available [20, 24]. As in the previous section, it is worth noticing the lack of a comparative study of the approaches. Again, the reason for this may be that the considered cross-domain task and data overlap scenarios vary among works.

5.3 Transferring Rating Patterns between Domains

Rather than sharing user or item latent factors for knowledge transfer, a different set of approaches analyzes the structure of rating data at the community level. These methods are based on the hypothesis that even when their users and items are different, close domains are likely to have user preferences sampled with the same population. Therefore, latent correlations may exist between preferences of groups of users for groups of items, which are referred to as *rating patterns*. In this context, rating patterns can act as a bridge that relates the domains (see Figure 10), such that knowledge transfer can be performed in either adaptive or collective manners.

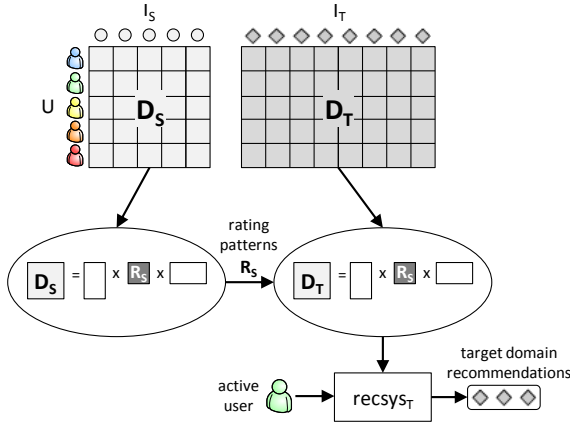


Fig. 10 *Transferring rating patterns.* A co-clustering model is learnt on the source domain to obtain rating patterns, which are used to cluster users and items in the target domain.

In the adaptive setting, rating patterns are extracted from a dense source domain. In the collective setting, data from all the domains are pulled together and jointly exploited, even though users and items do not overlap across domains.

Lee et al. [38] proposed one of the first approaches to exploit rating patterns for cross-domain recommendation. Similarly to the cross-domain mediation proposed by Berkovsky et al. [6], global nearest neighbors are identified by adding the similarity scores from each domain. Then, patterns of items commonly rated together by a set of neighbors are discovered using association rules. Finally, in the recommendation stage, rating predictions are computed with the standard user-based CF algorithm, but enhanced with the user’s rules that contain the target items.

Li et al. [40] proposed an adaptive method based on simultaneously co-clustering users and items in the source domain, to extract rating patterns. Clustering is performed using a tri-factorization of the source rating matrix [18]. Then, knowledge is transferred through a *codebook*, a compact cluster-level matrix computed in the source domain taking the average rating of each user-item cluster. In the target domain, missing ratings are predicted using the codebook. Moreno et al. [46] extended the codebook idea to a scenario in which various source domains contribute to the target domain. The approach is based on a linear combination of codebooks, where the weights are learnt by minimizing the prediction error in the target domain.

In a related work [41], Li et al. extended the same idea to a collective approach using a probabilistic framework. Instead of relying on an dense source domain data to build the codebook, all rating matrices are pulled together to extract the shared patterns. Furthermore, rather than having each user/item belonging to a single cluster, a probability distribution is introduced to allow users and items belong to multiple clusters, with distinct membership degrees. In the same fashion, the ratings associated with each user-item cluster are also given by a conditional probability distribution. In this way, a generative rating model is obtained, since the ratings of each

Table 8 Summary of cross-domain recommendation approaches based on transferring rating patterns between domains.

Cross-domain approach	Inter-domain relationships	References
<i>Extracting association rules from user rating behavior</i>	Rating overlap	Lee et al. 2001 [38] ^U
<i>Transferring implicit cluster-level rating patterns between domains</i>	Rating patterns	Li et al. 2009a [40] ^N Li et al. 2009b [41] ^N Moreno et al. 2012 [46] ^N Cremonesi & Quadrana 2014 [17] ^N
	Domain-independent parts of rating patterns	Gao et al. 2013 [26] ^N

(*N*) no overlap, (*U*) user overlap, (*I*) item overlap, (*UI*) user and item overlap.

domain can be recovered by drawing users and items from the shared cluster-level model, and then drawing the expected rating conditioned to the user-item cluster.

A strength of both approaches is that neither overlap of users nor of items is required. However, Cremonesi and Quadrana [17] partially disproved it, showing that the *codebook* does not transfer knowledge when source and target domains do not overlap. They provided an alternative explanation to the accuracy increase using a codebook that does not involve knowledge transfer between domains.

Finally, Gao et al. [26] followed the idea of extracting rating patterns by co-clustering rating matrices, and addressed two limitations of previous methods. First, they argued that some domains are more related to the target domain than others, and this cannot be captured using identical rating patterns. Second, they hypothesized that performance may suffer when the domains are diverse, and do not share common rating patterns. To overcome these limitations, the authors proposed a model capable of controlling the amount of knowledge transferred from each domain. Specifically, they used a co-clustering algorithm of [40], but split the extracted rating patterns into a shared part and a domain-specific part. In contrast to [40], optimization is performed in a collective way, since the shared part of the rating patterns is learnt simultaneously from all the domains.

Table 8 summarizes the described cross-domain approaches based on transferring rating patterns between domains. We observe that more recent methods based on clustering do not rely on any overlap between domains. However, as discussed in [26], care must be taken in order not to degrade performance by transferring noisy patterns from unrelated domains. We therefore conjecture that further research on the *when to transfer* aspect [49] will be conducted, to identify valuable information from source domains.




Fig. 11 Partitioning of \mathcal{D} : (left) *hold-out* – test ratings sampled and hidden without partitioning the users; (middle) *leave-some-users-out* – users split into disjoint training/test sets; (right) *leave-all-users-out* – ratings in the target dataset used as test profiles and ratings.

6 Evaluation of Cross-Domain Recommender Systems

In this section, we discuss the methods used to evaluate cross-domain recommender systems. The focal point is that such systems cannot be evaluated in a problem-independent way; whether a cross-domain recommender system is an appropriate solution cannot be evaluated without taking into account for what it is intended. The nature of the evaluation must be connected to the purpose for which the recommendations are required. Thus, we compare the corresponding evaluation methods based on the cross-domain recommendation goals addressed in the literature (see Section 2.3).

Three types of evaluations can be used to compare cross-domain recommender systems [25, 57]. *Offline experiments* evaluate a system by analyzing past user preferences. They are typically the easiest to conduct, as they require no interaction with real users. With *online studies*, a small group of subjects is asked to use the system in a controlled environment, and to report on the experience. Finally, *live trials* evaluates the system based on feedback from real users. As most cross-domain recommendation works use offline experiments (with a few performing online studies, and no live trials, see Table 9), we focus on offline experiments. The reader is referred to Chapter ?? for an extensive discussion on methodologies and metrics used to evaluate recommender systems.

The decision regarding the evaluation method is often critical, as each one reflects a specific task or goal. Many offline evaluation schemes exist, which differ in a number of aspects: *data partitioning*, *metrics*, and *sensitivity analysis* (e.g., relative density of domain datasets, and degree of overlap between domains), as discussed respectively in the next sections.

6.1 Data Partitioning

In order to evaluate algorithms offline, it is necessary to simulate the process where the system makes recommendations, and users evaluate them. This requires pre-recorded datasets of interactions between users and items. In cross-domain applica-

Table 9 Summary of cross-domain recommendation approaches based on the technique used to partition the data into training and test sets.

Data partitioning	References	
<i>Online studies</i>	Braunhofer et al. 2013 [9] Fernandez-Tobias et al. 2013 [23] Shapira et al. 2013 [58]	Szomszor et al. 2008 [61] Winoto et al. 2008 [66]
<i>Leave-all-users-out</i>	Cremonesi et al. 2011 [16] Goga et al. 2013 [28] Hu et al. 2013 [31] Jain et al. 2013 [32]	Kaminskas et al. 2013 [35] Loni et al. 2014 [44] Shapira et al. 2013 [58] Tiroshi et al. 2012 [65]
<i>Leave-some-users-out</i>	Abel et al. 2011 [1] Abel et al. 2013 [2]	Li et al. 2009 [40, 41] Stewart et al. 2009 [60]
<i>Hold-out</i>	Li et al. 2011 [42] Nakatsuji et al. 2010 [47] Pan et al. 2008 [48] Pan et al. 2010 [51] Pan et al. 2012 [52] Pan et al. 2013 [53]	Sahebi et al. 2013 [55] Shi et al. 2011 [59] Tang et al. 2011 [63] Zhang et al. 2010 [67] Zhang et al. 2013 [68] Zhao et al. 2013 [69]

tions, there are (at least) two potentially overlapping datasets: the source dataset \mathcal{D}_S and the target dataset \mathcal{D}_T .

We assume \mathcal{D}_S and \mathcal{D}_T are chosen according to the recommendation task and goal in hand. For instance, if we are evaluating a cross-selling recommender, \mathcal{D}_S and \mathcal{D}_T are set at the *item level* as described in Section 2.1, contain items of different nature, like movies and books, and have overlapping users. On the contrary, if we are evaluating a cross-domain recommender as a tool to increase recommendation diversity, \mathcal{D}_S and \mathcal{D}_T are set at the *item attribute level*, with items of the same type, but differ in the value of certain attribute, as comedy and drama movies.

In offline evaluations, a portion of \mathcal{D}_T is hidden to facilitate prediction of the available knowledge, and gauge the quality of the recommendations. There is a number of ways to choose the ratings to be hidden. The most general approach creates three subsets of ratings from the original datasets: (i) $\mathcal{D}_{\text{training_profiles}}$, which contains the set of ratings from users $\mathcal{U}_{\text{training_profiles}}$ for items $\mathcal{I}_{\text{training_profiles}}$ that are used to train the algorithms under evaluation; (ii) $\mathcal{D}_{\text{test_profiles}}$, which contains the set of users $\mathcal{U}_{\text{test_profiles}}$ and their known ratings for items $\mathcal{I}_{\text{test_profiles}}$ that are used as input profiles for the trained recommender; and (iii) $\mathcal{D}_{\text{test_ratings}}$, which contains the set of users $\mathcal{U}_{\text{test_profiles}}$ and their hidden ratings for items $\mathcal{I}_{\text{test_ratings}}$ that are used as the ground truth to evaluate the recommendations.

Depending on the choice of the $\mathcal{D}_{\text{training_profiles}}$, $\mathcal{D}_{\text{test_profiles}}$, and $\mathcal{D}_{\text{test_ratings}}$ subsets, different evaluation data partitions can be designed.

- *Hold-out* (Figure 11-left) is implemented when $\mathcal{D}_{\text{test_profiles}} \subseteq \mathcal{D}_{\text{training_profiles}}$, i.e., test ratings are sampled and hidden from the original dataset without partitioning the users. This partition is suitable to evaluate linked- and multi-domain rec-

ommenders with the accuracy goal, and is applicable to memory-based recommenders, which are unable to provide recommendations to new users.

- *Leave-some-users-out* (Figure 11-middle) is implemented when $\mathcal{U}_{\text{training_profiles}} \cap \mathcal{U}_{\text{test_profiles}} = \emptyset$, i.e., the users are split into two disjoint subsets: one for training and one for testing. This partition is suitable to evaluate a cross-domain recommender with the new user goal.
- *Leave-all-users-out* (Figure 11-right) is implemented when $\mathcal{D}_{\text{training_profiles}} \cap \mathcal{D}_T = \emptyset$, i.e., the ratings in the target dataset are used only as profile and test ratings. This partition is suitable to evaluate a cross-domain recommender with the cold-start and new item goals.

6.2 Metrics

The notion of relevance of recommendations and the ways to measure it have been debated in numerous works on recommender systems. Generally speaking, there are three categories of evaluation metrics: *predictive metrics*, *ranking metrics*, and *classification metrics* [30].

Theoretical debates surrounds the distribution of the missing ratings. Because of the data sparsity, offline evaluations are performed on a small fraction of the available items. Each metric makes implicit assumptions regarding the value and the distribution of the missing ratings, which impact the interpretation of obtained results. For instance, predictive metrics like *MAE* and *RMSE* assume that the unknown ratings are missing at random, the classification metric of *precision* assumes that all missing ratings are irrelevant for the user, whereas *recall*, *fallout*, and *ROC* assume that non-relevant ratings are missing with a higher probability than relevant ratings. Practical debates also consider the recommendation goal. Prediction metrics are to be preferred when the goal is to reduce the sparsity of the target domain; ranking metrics are adopted when testing user models, especially in cold-start situations; and classification metrics are best-suited for the top-N recommendation task.

Table 10 summarizes the offline evaluation metrics exploited in cross-domain recommenders. The majority of works adopts prediction metrics. This is motivated by the fact that the addressed goal is to reduce sparsity and increase accuracy, and the algorithms designed for this are often based on error-metric optimization techniques, which are naturally evaluated using the category of predictive metrics.

6.3 Sensitivity Analysis

The performance of a cross-domain recommender is mainly affected by three parameters: the overlap between the source and target domains, the density of the target domain data, and the size of the target user's profile. Thus, the evaluation of a cross-domain recommender system mostly considered the sensitivity of the algorithms with respect to these three parameters.

Table 10 Summary of metrics used for the evaluation of cross-domain recommender system.

Category	Metric	References	
<i>Prediction metrics</i>	<i>MAE</i>	Berkovsky et al. 2007 [6, 7]	Pan et al. 2008 [48]
		Berkovsky et al. 2008 [8]	Pan et al. 2010 [51]
		Cao et al. 2010 [13]	Pan et al. 2012 [52]
		Hu et al. 2013 [31]	Pan et al. 2013 [53]
		Li et al. 2009 [40, 41]	Shapira et al. 2013 [58]
		Moreno et al. 2012 [46]	Shi et al. 2011 [59]
		Loni et al. 2014 [44]	Winoto et al. 2008 [66]
		Nakatsuji et al. 2010 [47]	
	<i>RMSE</i>	Li et al. 2011 [42]	Pan et al. 2013 [53]
		Loni et al. 2014 [44]	Sahebi et al. 2013 [55]
		Pan et al. 2010 [51]	Zhang et al. 2010 [67]
		Pan et al. 2012 [52]	Zhao et al. 2013 [69]
<i>Ranking metrics</i>	<i>ROC</i>	Goga et al. 2013 [28]	
	<i>MRR</i>	Abel et al. 2011 [1]	Abel et al. 2013 [2]
	<i>nDCG</i>	Zhang et al. 2013 [68]	
	<i>AUC</i>	Fernandez-Tobias et al. 2013 [23] Hu et al. 2013 [31]	Tiroshi et al. 2012 [65]
	<i>MAP</i>	Fernández-Tobías et al. 2013 [23] Shapira et al. 2013 [58] Jain et al. 2013 [32]	Shapira et al. 2013 [58] Zhang et al. 2013 [68]
<i>Classification metrics</i>	<i>Precision</i>	Kaminskas et al. 2013 [35] Tiroshi et al. 2013 [64]	Stewart et al. 2009 [60]
	<i>Recall</i>	Stewart et al. 2009 [60]	Nakatsuji et al. 2010 [47]
	<i>F-measure</i>	Cremonesi et al. 2011 [16]	Gao et al. 2013 [26]

Most works have assumed an overlap of users between the source and target domains. They all conducted evaluations with 100% of overlap, except for two works. Cremonesi et al. [16] analyzed the behavior of various cross-domain recommenders by varying the percentage of user-overlap in the range 0%-50%, and Zhao et al. [69] adopted a similar evaluation by varying the percentage of user overlap in the range 0%-100%. Fewer works [8, 16, 53, 69] studied the case of item overlap, and they all assume to have the same catalog of items across domains. Some works [2, 9, 23, 35, 60, 61] studied the case of overlapping features, especially social tags. Shi et al. [59] studied the sensitivity of the cross-domain recommender by varying the number of overlapping tags between 5 and 50.

Some works [8, 40, 41, 55, 59] have studied the sensitivity of recommendations as a function of the user profile size, i.e., the number of ratings provided by the user receiving the recommendations. This is particularly important for the cold-start and new user goals. Both Pan et al. [51] and Abel et al. [2] developed tag-based recommenders, and performed their analysis by varying the number of tags in the

Table 11 Summary of variables for sensitivity analysis of cross-domain recommender systems.

Parameter	References	
<i>Overlap between domains</i>	Abel et al. 2013 [2] Cremonesi et al. 2011 [16]	Shi et al. 2011 [59] Zhao et al. 2013 [69]
<i>Target domain density</i>	Cao et al. 2010 [13] Cremonesi et al. 2011 [16] Pan et al. 2008 [48]	Pan et al. 2010 [51] Shapira et al. 2013 [58]
<i>User profile size</i>	Berkovsky et al. [6, 7] Berkovsky et al. 2008 [8] Li et al 2009 [40, 41]	Sahebi et al. 2013 [55] Shi et al. 2011 [59]

user profile in the 10–40 and 0–150 ranges, respectively. Others conducted a similar analysis on rating-based recommenders: Shi et al. 2011 [59] varied the profile size from 20 to 100 ratings, Berkovsky et al. [8] varied the profile size from 3% to 33% of ratings, and Li et al. [40, 41] and Sahebi et al. [55] varied the profile size in the range of 5–15 and 1–20 ratings, respectively.

Finally, some works [13, 16, 51, 58] have studied the quality of recommendations as a function of the dataset density. This is important for the cold-start and accuracy goals. Cao et al. [13] varied the density of the multi-domain dataset, i.e., the union of source and target datasets, between 0.2% and 1%. Shapira et al. [58] varied the density of the dataset between 1% and 40%, but only for the baseline single-domain algorithms, while evaluating cross-domain algorithms at the 1% density. Cremonesi et al. [16] varied the density of the target domain between 0.1% and 0.9%. The sensitivity analyses performed in the above works are summarized in Table 11.

7 Practical Considerations in Cross-Domain Recommendation

We have covered so far a wide spectrum of models and techniques applicable to cross-domain recommendation. Recommender system practitioners may find this variety of options overwhelming, when materializing a cross-domain recommender. Therefore, we list several practical considerations that drive the choice of the appropriate recommendation solution.

The first set of considerations deals with the pivotal questions of “*what, when, and how to transfer?*” that have already been raised in Section 3. The term ‘transfer’ refers in the following discussion to both the knowledge aggregation (Section 4) and the knowledge transfer (Section 5) approaches.

- *What to transfer?* Single-domain recommenders may gather different types of user data: explicit ratings, unary purchase lists, browsing logs, and many others. They are also likely to store domain metadata and recommendation method-specific data, e.g., collaborative neighborhoods of similar users and matrix factor-

ization latent vectors. It cannot be determined in advance what knowledge from the source domain recommenders can benefit the target domain recommender, and some form of information gain analysis needs to be done. This may be a complex process, in which the target recommendation method and the recommendation task in hand should be taken into consideration .

- *When to transfer?* Deciding what information should be transferred is tightly bound to the consideration of conditions under which the transfer is beneficial. It is clear that at the initial deployment period of the target domain recommender, the transfer will enrich the recommender. On the contrary, no transfer is needed when the target domain recommender possesses complete and up-to-date information. But what happens in-between? This depends not only on the sparsity of the target domain data, but also on factors like the overlap of ratings between domains, and freshness of data in the source domains.
- *How to transfer?* The answer to the ‘how’ question deals with the implementation of the knowledge transfer. Two high-level options are possible: either to implement direct one-to-one mappings between the source and the target recommenders, or to leverage a common representation that will facilitate the transfer. The downside of the former is that the number of possible combinations is quadratic and will grow if new recommenders are being introduced. The latter requires only a single transfer mechanism from/to the common representation, but an agreed upon representation is hard to achieve in practice. Some rules for reconciling conflicts in the transferred data should also be put in place.

Additional question that needs to be dealt with is “*where from to transfer?*” This question is peripheral in transfer learning since any available information is considered relevant, but this is not the case in cross-domain recommenders. The main indicator here is the the distance between domains. Some pairs of domains, e.g., movies and TV, are inherently closer than others, e.g., games and tourism. The close domains have a greater potential to benefit the target recommender, and are naturally the preferred sources. Contextual factors (location, temporal closeness) and the overlaps of user and item sets are also important in answering this question. We believe that practical cross-domain recommenders need to thoroughly examine the sources of the transferred knowledge.

Knowledge transfer between domains typically requires some *auxiliary information*. We highlight here two types of such information, which actually underpin the transfer. These are semantic networks like WordNet and DBpedia, and open or crowdsourced knowledge references like Wikipedia and Open Directory. The auxiliary information is critical for the knowledge transfer, since it links the domains and informs the answer to the ‘how’ question. Hence, important considerations faced by a practical cross-domain recommender deal with the the availability and reliability of the auxiliary information. Chapters ?? and ?? address such issues in semantic-aware and social recommender systems.

The next set of considerations deals with the target *recommendation task*. Many options exist here: best item vs. top-K, one-off vs. sequential interaction, single product vs. bundle of products, recommendation to individual users vs. to a group of a users. Every recommendation scenario implies a different algorithm in place, and

also distinct types of knowledge that can be transferred from the source domains. Related to this, the *metric* of recommendation success should be considered. Do the recommendations need to discover all the relevant items, match as many aspects of user interests as possible, or provide a surprising recommendation? Likewise, *technical constraints* may be an important factor. For instance, are the recommendations computed offline or delivered live to users? Is it a server-side recommendation which can be resource intense, or a lightweight client-side recommendation? These considerations cannot be discarded, as answers to the above questions may affect the choice of the knowledge transfer and of the cross-domain recommendation approach.

Last but not the least, special attention should be paid to *ethical* and *privacy aspects* (Chapter ??) in cross-domain recommenders. Transferring data and knowledge between single-domain recommenders may contradict privacy policies of the recommenders and existing privacy regulations. Moreover, it may allow malicious attackers not only to get access to a larger volume of user data, but also to apply data mining to the combined knowledge, uncovering (potentially sensitive) information. With respect to this, knowledge transfer methods are generally more robust than the aggregation methods, although they still cannot completely eliminate the data mining risk. Developers of a cross-domain recommender should keep the privacy consideration in mind, when selecting their knowledge transfer method.

8 Open Research Issues

This section provides an overview of new requirements and applications emerging from the landscape of cross-domain recommender systems. One interesting issue that deserves more attention in the future is the synergy between contextual and cross-domain recommendations: different contexts (e.g., location, time, and mood) can be treated as different domains (see Chapter ?? for details on context-aware recommendation). This opens interesting scenarios in which context-aware techniques can be applied to cross-domain recommendations, and vice versa. Moreover, context can be treated as a bridge between different domains, and seminal work has already been carried out in this direction [9, 23].

Another important issue concerns the metrics adopted for the evaluation of the recommendations. A common practice with cross-domain recommender systems is to evaluate their relevance through predictive accuracy metrics, such as MAE and RMSE, which capture the error between the actual and predicted ratings. However, in many commercial systems only a small number of best recommendations is shown, while the predicted ratings are not. That is, the system suggests a few items that are likely to be very appealing for users. Direct evaluation of top-N recommendation performance must be accomplished by means of alternative methodologies based either on classification metrics (e.g., recall and fallout) or ranking metrics (e.g., average reciprocal hit-rank and average relative position), as explained in Chapter ??.

We can push this idea further, by considering that accuracy is not sufficient to provide useful recommendations. Other criteria have been proposed to augment the evaluation dimensions, such as diversity, novelty, and serendipity (see Chapter ??). As one can expect, cross-domain recommendations would be less accurate than those based on the same amount of user data pertaining to the target domain. However, the true advantage of cross-domain recommendations is not necessarily in their accuracy, but rather in their novelty and diversity, which may lead to a higher satisfaction and utility for the user. In this context, the recently proposed novelty and diversity metrics could be taken into consideration [57].

The next open research issue deals with the use of cross-domain recommender systems as a means to reduce the user model elicitation effort. The preference elicitation process is important for the recommenders (Chapter ??), but it may pose two conflicting requirements. On the one hand, the system must collect “enough” ratings in order to learn the users’ preferences and improve the accuracy of recommendations. On the other hand, gathering ratings imposes a burden on the users, which may negatively affect their experience. Cross-domain recommender systems could be used as alternative elicitation tools able to build detailed user profiles without the need to collect explicit user preferences.

Finally, the importance of real life datasets needs to be stressed (Chapter ??). These are necessary for evaluations of new cross-domain approaches, but are quite scarce and hard to reach in practice. Large-scale cross-domain datasets are gathered by big industry players, like Amazon, eBay, and Yelp, but these datasets rarely become available to the broader research community. We would like to encourage industry researchers to cooperate with the academic researchers and share their data. This could boost both the research in cross-domain recommendation and the deployment of practical cross-domain recommenders.

References

1. Abel, F., Araújo, S., Gao, Q., Houben, G.-J.: Analyzing Cross-system User Modeling on the Social Web. *11th International Conference on Web Engineering*, pp. 28–43 (2011)
2. Abel, F., Helder, E., Houben, G.-J., Henze, N., Krause, D.: Cross-system User Modeling and Personalization on the Social Web. *User Modeling and User-Adapted Interaction* 23(2-3), pp. 169–209 (2013)
3. Azak, M.: Crossing: A Framework to Develop Knowledge-based Recommenders in Cross Domains. *MSc thesis, Middle East Technical University* (2010)
4. Berkovsky, S., Kuflik, T., Ricci, F.: Entertainment Personalization Mechanism through Cross-domain User Modeling. *1st International Conference on Intelligent Technologies for Interactive Entertainment*, pp. 215–219 (2005)
5. Berkovsky, S., Goldwasser, D., Kuflik, T., Ricci, F.: Identifying Inter-domain Similarities through Content-based Analysis of Hierarchical Web-Directories. *17th European Conference on Artificial Intelligence*, pp. 789–790 (2006)
6. Berkovsky, S., Kuflik, T., Ricci, F.: Cross-Domain Mediation in Collaborative Filtering. *11th International Conference on User Modeling*, pp. 355–359 (2007)
7. Berkovsky, S., Kuflik, T., Ricci, F.: Distributed Collaborative Filtering with Domain Specialization. *1st ACM Conference on Recommender Systems*, pp. 33–40 (2007)

8. Berkovsky, S., Kuflik, T., Ricci, F.: Mediation of User Models for Enhanced Personalization in Recommender Systems. *User Modeling and User-Adapted Interaction* 18(3), pp. 245–286 (2008)
9. Braunhofer, M., Kaminskas, M., Ricci, F.: Location-aware Music Recommendation. *International Journal of Multimedia Information Retrieval* 2(1), pp. 31–44 (2013)
10. Cantador, I., Fernández-Tobías, I., Bellogín, A., Kosinski, M., Stillwell, D.: Relating Personality Types with User Preferences in Multiple Entertainment Domains. *1st Workshop on Emotions and Personality in Personalized Services*, CEUR workshop Proceedings, vol. 997 (2013)
11. Carmagnola, F., Cena, F.: User Identification for Cross-system Personalisation. *Information Sciences* 179(1-2), pp. 16–32 (2009)
12. Carmagnola, F., Cena, F., Gena, C.: User Model Interoperability: A Survey. *User Modeling and User-Adapted Interaction* 21(3), pp. 285–331 (2011)
13. Cao, B., Liu, N.N., Yang, Q.: Transfer Learning for Collective Link Prediction in Multiple Heterogeneous Domains. *27th International Conference on Machine Learning*, pp. 159–166 (2010)
14. Chung, R., Sundaram, D., Srinivasan, A.: 2007. Integrated Personal Recommender Systems. *9th International Conference on Electronic Commerce*, pp. 65–74 (2007)
15. Costa, P. T., McCrae, R. R.: Revised NEO Personality Inventory (NEO-PI-R) and NEO Five-Factor Inventory (NEO-FFI) Manual. *Psychological Assessment Resources* (1992)
16. Cremonesi, P., Tripodi, A., Turrin, R.: Cross-domain Recommender Systems. *11th IEEE International Conference on Data Mining Workshops*, pp. 496–503 (2011)
17. Cremonesi, P., Quadrana, M.: Cross-domain recommendations without overlapping data: myth or reality? *8th ACM Conference on Recommender Systems* (2014)
18. Ding, C., Li, T., Peng, W., Park, H.: Orthogonal Nonnegative Matrix Tri-factorizations for Clustering. *12th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 126–135 (2006)
19. Driskill, R., Riedl, J.: Recommender Systems for E-Commerce: Challenges and Opportunities. *AAAI'99 Workshop on Artificial Intelligence for Electronic Commerce*, pp. 73–76 (1999)
20. Enrich, M., Braunhofer, M., Ricci, F.: Cold-Start Management with Cross-Domain Collaborative Filtering and Tags. *14th International Conference on E-Commerce and Web Technologies*, pp. 101–112 (2013)
21. Fernández-Tobías, I., Cantador, I., Kaminskas, M., Ricci, F.: 2011. A Generic Semantic-based Framework for Cross-domain Recommendation. *2nd International Workshop on Information Heterogeneity and Fusion in Recommender Systems*, pp. 25–32 (2011)
22. Fernández-Tobías, I., Cantador, I., Kaminskas, M., Ricci, F.: Cross-domain Recommender Systems: A Survey of the State of the Art. *2nd Spanish Conference on Information Retrieval*, pp. 187–198 (2012)
23. Fernández-Tobías, I., Cantador, I., Plaza, L.: An Emotion Dimensional Model Based on Social Tags: Crossing Folksonomies and Enhancing Recommendations. *14th International Conference on E-Commerce and Web Technologies*, pp. 88–100 (2013)
24. Fernández-Tobías, I., Cantador, I.: Exploiting Social Tags in Matrix Factorization Models for Cross-domain Collaborative Filtering. *1st International Workshop on New Trends in Content-based Recommender Systems* (2013)
25. Freyne, J., Berkovsky, S., Smith, G.: Evaluating Recommender Systems for Supportive Technologies. *User Modeling and Adaptation for Daily Routines*, pp. 195–217 (2013)
26. Gao, S., Luo, H., Chen, D., Li, S., Gallinari, P., Guo, J.: Cross-Domain Recommendation via Cluster-Level Latent Factor Model. *17th and 24th European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 161–176 (2013)
27. Givon, S., Lavrenko, V.: Predicting Social-tags for Cold Start Book Recommendations. *3rd ACM Conference on Recommender Systems*, pp. 333–336 (2009)
28. Goga, O., Lei, H., Parthasarathi, S. H. K., Friedland, G., Sommer, R., Teixeira, R.: Exploiting Innocuous Activity for Correlating Users across Sites. *22nd International Conference on World Wide Web*, pp. 447–458 (2013)

29. González, G., López, B., de la Rosa, J. LL.: A Multi-agent Smart User Model for Cross-domain Recommender Systems. In: *Beyond Personalization 2005 - The Next Stage of Recommender Systems Research*, pp. 93–94 (2005)
30. Hellocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.: Evaluating Collaborative Filtering Recommender Systems. *ACM Transactions on Information Systems* 22(1), pp. 5–53 (2004)
31. Hu, L., Cao, J., Xu, G., Cao, L., Gu, Z., Zhu, C.: Personalized Recommendation via Cross-domain Triadic Factorization. *22nd International Conference on World Wide Web*, pp. 595–606 (2013)
32. Jain, P., Kumaraguru, P., Joshi, A.: @i seek 'fb.me': Identifying Users across Multiple Online Social Networks. *22nd International Conference on WWW Companion*, pp. 1259–1268 (2013)
33. Jialin Pan, S., Yang, Q.: A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering* 22(10), pp. 1345–1359 (2010)
34. Joon Kook, H.: Profiling Multiple Domains of User Interests and Using them for Personalized Web Support. *1st International Conference on Intelligent Computing*, pp. 512–520 (2005)
35. Kaminskas, M., Fernández-Tobías, I., Ricci, F., Cantador, I.: Ontology-based Identification of Music for Places. *13th International Conference on Information and Communication Technologies in Tourism*, pp. 436–447 (2013)
36. Kitts, B., Freed, D., Vrieze, M.: Cross-sell: A Fast Promotion-tunable Customer-item Recommendation Method based on Conditionally Independent Probabilities. *6th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 437–446 (2000)
37. Koren, Y.: Factorization Meets the Neighborhood: A Multifaceted Collaborative Filtering Model. *14th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 426–434 (2008)
38. Lee, C. H., Kim, Y. H., Rhee, P. K.: Web Personalization Expert with Combining Collaborative Filtering and Association Rule Mining Technique. *Expert Systems with Applications* 21(3), pp. 131–137 (2001)
39. Li, B.: Cross-Domain Collaborative Filtering: A Brief Survey. *23rd IEEE International Conference on Tools with Artificial Intelligence*, pp. 1085–1086 (2011)
40. Li, B., Yang, Q., Xue, X.: Can Movies and Books Collaborate? Cross-domain Collaborative Filtering for Sparsity Reduction. *21st International Joint Conference on Artificial Intelligence*, pp. 2052–2057 (2009)
41. Li, B., Yang, Q., Xue, X.: Transfer Learning for Collaborative Filtering via a Rating-matrix Generative Model. *26th International Conference on Machine Learning*, pp. 617–624 (2009)
42. Li, B., Zhu, X., Li, R., Zhang, C., Xue, X., Wu, X.: Cross-domain Collaborative Filtering over Time. *22nd International Joint Conference on Artificial Intelligence*, pp. 2293–2298 (2011)
43. Loizou, A.: How to Recommend Music to Film Buffs: Enabling the Provision of Recommendations from Multiple Domains. *PhD thesis, University of Southampton* (2009)
44. Loni, B, Shi, Y, Larson, M. A., Hanjalic, A.: Cross-Domain Collaborative Filtering with Factorization Machines. *36th European Conference on Information Retrieval* (2014)
45. Low, Y., Agarwal, D., Smola, A. J.: Multiple Domain User Personalization. *17th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 123–131 (2011)
46. Moreno, O. Shapira, B. Rokach, L. Shani, G.: TALMUD: transfer learning for multiple domains. *21st ACM Conference on Information and Knowledge Management*, pp. 425–434 (2012)
47. Nakatsuji, M., Fujiwara, Y., Tanaka, A., Uchiyama, T., Ishida, T.: Recommendations Over Domain Specific User Graphs. *19th European Conference on Artificial Intelligence*, pp. 607–612 (2010)
48. Pan, S. J., Kwok, J. T., Yang, Q.: Transfer Learning via Dimensionality Reduction. *23rd AAAI Conference on Artificial Intelligence*, pp. 677–682 (2008)
49. Pan, S. J., Yang, Q.: A Survey on Transfer Learning. *IEEE Transactions on Knowledge and Data Engineering* 22(10), pp. 1345–1359 (2010)
50. Pan, W., Liu, N. N., Xiang, E. W., Yang, Q.: Transfer Learning to Predict Missing Ratings via Heterogeneous User Feedbacks. *22nd International Joint Conference on Artificial Intelligence*, pp. 2318–2323 (2011)

51. Pan, W., Xiang, E.W., Liu, N.N., Yang, Q.: Transfer Learning in Collaborative Filtering for Sparsity Reduction. *24th AAAI Conference on Artificial Intelligence*, pp.210–235 (2010)
52. Pan, W., Xiang, E. W., Yang, Q.: Transfer Learning in Collaborative Filtering with Uncertain Ratings. *26th AAAI Conference on Artificial Intelligence*, pp. 662–668 (2012)
53. Pan, W., Yang, Q.: Transfer Learning in Heterogeneous Collaborative Filtering Domains. *Artificial Intelligence* 197, pp. 39–55 (2013)
54. Rendle, S.: Factorization Machines with libFM. *ACM Transactions on Intelligent Systems and Technology* 3(3), pp. 1–22 (2012)
55. Sahebi, S., Brusilovsky, P.: Cross-Domain Collaborative Recommendation in a Cold-Start Context: The Impact of User Profile Size on the Quality of Recommendation. *21st International Conference on User Modeling, Adaptation, and Personalization*, pp. 289–295 (2013)
56. Salakhutdinov, R., Mnih, A.: Probabilistic Matrix Factorization. *Advances in Neural Information Processing Systems* 20, pp. 1257–1264 (2008)
57. Shani, G., Gunawardana, A.: Evaluating Recommendation Systems. *Recommender Systems Handbook*, pp. 257–297 (2011)
58. Shapira, B., Rokach, L., Freilikhman, S.: Facebook Single and Cross Domain Data for Recommendation Systems. *User Modeling and User-Adapted Interaction* 23(2-3), pp. 211–247 (2013)
59. Shi, Y., Larson, M., Hanjalic, A.: Tags as Bridges between Domains: Improving Recommendation with Tag-induced Cross-domain Collaborative Filtering. *19th International Conference on User Modeling, Adaptation, and Personalization*, pp. 305–316 (2011)
60. Stewart, A., Diaz-Aviles, E., Nejdl, W., Marinho, L. B., Nanopoulos, A., Schmidt-Thieme, L.: Cross-tagging for Personalized Open Social Networking. *20th ACM Conference on Hypertext and Hypermedia*, pp. 271–278 (2009)
61. Szomszor, M. N., Alani, H., Cantador, I., O’Hara, K., Shadbolt, N.: Semantic Modelling of User Interests Based on Cross-Folksonomy Analysis. *7th International Semantic Web Conference*, pp. 632–648 (2008)
62. Szomszor, M. N., Cantador, I., Alani, H.: Correlating User Profiles from Multiple Folksonomies. *19th ACM Conference on Hypertext and Hypermedia*, pp. 33–42 (2008)
63. Tang, J., Yan, J., Ji, L., Zhang, M., Guo, S., Liu, N., Wang, X., Chen, Z.: Collaborative Users’ Brand Preference Mining across Multiple Domains from Implicit Feedbacks. *25th AAAI Conference on Artificial Intelligence*, pp. 477–482 (2011)
64. Tiroshi, A., Berkovsky, S., Kaafar, M. A., Chen, T., Kufflik, T.: Cross Social Networks Interests Predictions Based on Graph Features. *7th ACM Conference on Recommender Systems*, pp. 319–322 (2013)
65. Tiroshi, A., Kufflik, T.: Domain Ranking for Cross Domain Collaborative Filtering. *20th International Conference on User Modeling, Adaptation, and Personalization*, pp. 328–333 (2012)
66. Winoto, P., Tang, T.: If You Like the Devil Wears Prada the Book, Will You also Enjoy the Devil Wears Prada the Movie? A Study of Cross-Domain Recommendations. *New Generation Computing* 26, pp. 209–225 (2008)
67. Zhang, Y., Cao, B., Yeung, D.-Y.: Multi-Domain Collaborative Filtering. *26th Conference on Uncertainty in Artificial Intelligence*, pp. 725–732 (2010)
68. Zhang, X., Cheng, J., Yuan, T., Niu, B., Lu, H.: TopRec: Domain-specific Recommendation through Community Topic Mining in Social Network. *22nd International Conference on World Wide Web*, pp. 1501–1510 (2013)
69. Zhao, L., Pan, S. J., Xiang, E. W., Zhong, E., Lu, X., Yang, Q.: Active Transfer Learning for Cross-System Recommendation. *27th AAAI Conference on Artificial Intelligence*, pp. 1205–1211 (2013)
70. Zhuang, F., Luo, P., Xiong, H., Xiong, Y., He, Q., Shi, Z.: Cross-domain Learning from Multiple Sources: A Consensus Regularization Perspective. *IEEE Transactions on Knowledge and Data Engineering* 22(12), pp. 1664–1678 (2010)