ELSEVIER



Online Social Networks and Media



journal homepage: www.elsevier.com/locate/osnem

Recommender Systems for Online and Mobile Social Networks: A survey

Mattia G. Campana*, Franca Delmastro

IIT-CNR, Via G. Moruzzi 1, 56124 Pisa, Italy

ARTICLE INFO

Article history: Received 24 September 2017 Revised 28 October 2017 Accepted 30 October 2017

Keywords: Recommender Systems Online Social Networks Mobile Social Networks

ABSTRACT

Recommender Systems (RS) currently represent a fundamental tool in online services, especially with the advent of Online Social Networks (OSN). In this case, users generate huge amounts of contents and they can be quickly overloaded by useless information. At the same time, social media represent an important source of information to characterize contents and users' interests. RS can exploit this information to further personalize suggestions and improve the recommendation process. In this paper we present a survey of Recommender Systems designed and implemented for Online and Mobile Social Networks, highlighting how the use of social context information improves the recommendation task, and how standard algorithms must be enhanced and optimized to run in a fully distributed environment, as opportunistic networks. We describe advantages and drawbacks of these systems in terms of algorithms, target domains, evaluation metrics and performance evaluations. Eventually, we present some open research challenges in this area.

© 2017 Elsevier B.V. All rights reserved.

Contents

1.	Introduction	76
2.	The recommendation task	76
	2.1. Evaluation metrics	77
	2.1.1. Prediction accuracy	77
	2.1.2. Ranking accuracy.	77
3.	Main approaches for Recommender Systems	78
	3.1. Collaborative Filtering	78
	3.1.1. Memory-based	78
	3.1.2. Model-based	78
	3.2. Content-based Recommender Systems	79
	3.3. Graph-based	79
	3.4. Context-Aware Recommender Systems	80
4.	Recommender systems for Online Social Networks	83
	4.1. Social-aware recommendations	84
	4.1.1. Friends and people recommendations	84
	4.1.2. Items recommendations	85
	4.2. Tag-based recommendations	86
	4.3. Location-aware recommendations	87
5.	Recommender Systems for Mobile Social Networks	89
	5.1. Collaborative Filtering	89
	5.1.1. OppCF	89
	5.1.2. MobHinter	90

* Corresponding author.

E-mail addresses: m.campana@iit.cnr.it (M.G. Campana), f.delmastro@iit.cnr.it (F. Delmastro).

https://doi.org/10.1016/j.osnem.2017.10.005 2468-6964/© 2017 Elsevier B.V. All rights reserved.

5.1.3. DiffeRS	90							
5.1.4. LocPref	90							
5.2. Exploiting tags	91							
5.2.1. Information-Centric Habit (ICe-Habit)	91							
5.2.2. Pervasive PLIERS (p-PLIERS)	91							
5. Concluding remarks and open challenges	92							
Acknowledgments								
References								

1. Introduction

Recommender Systems (RS) and Online Social Networks (OSN) have established a strong cooperation in the last few years. They both aim at coping with the huge amount of data produced and shared by users through online platforms, trying to maintain a high user engagement. This cooperation is built upon the advantages that both systems can achieve: the optimization of the recommendation techniques, by exploiting additional content and user characterization derived from OSN, and the increasing request of OSN services' personalization. RS targeting the social media domain have been already defined in the literature as Social Recommender Systems (SRS) [1,2]. However, this notion should be further expanded, embracing also the recent evolution of social media towards the pervasive and mobile computing environment. Here, the network can be also implemented by the physical co-location of users and devices, and their ability to exploit direct wireless communications for content sharing and dissemination, without requiring a constant Internet access. In this direction, the concept of Mobile Social Networks (MSN) has been introduced in [3], not as a simple extension of OSN services running on mobile devices, but as the opportunity to create real networks of people, in which users actively participate in the generation and sharing of contents, anywhere and anytime, based on opportunistic networking and device-to-device (D2D) communications. In this scenario, RS can be optimized by exploiting additional information derived from users' mobile devices (i.e., sensors), which help contextualize the user's preferences, and to rely on a partial and dynamic knowledge of the network and the available data. This concept contributes to the migration towards a new Internet paradigm, the Internet of People (IoP) [4], in which users are at the center, by actively contributing to the network evolution both from the communication and data point of views.

In this paper we present a survey of RS defined both for OSN and MSN, highlighting advantages and drawbacks of standard recommendation techniques applied in these environments, and how those systems evolved over time, in terms of technical solutions, target domains, evaluation metrics and performance evaluations. In Section 2, we summarize the problem of recommendations and the evaluation metrics generally used in the literature. Then, in Section 3, we present an overview of standard techniques like Collaborative Filtering (CF), Content-based RS, Network-based RS and Context-aware RS, with particular attention to the various information used to characterize the relationship between users and items to further optimize and personalize the recommendations. This section would also introduce the reader with standard notations and methods applied in RS, which will be analyzed in the subsequent sections. In Sections 4 and 5, we describe the main solutions presented in the literature for OSN and MSN, respectively. The first area has been widely studied in the last years, proposing solutions that can address different target domains (e.g., to recommend people, locations, point of interests, tags or contents) by exploiting heterogeneous context information. Therefore, we decided to group the proposed solutions by the type of context information used (e.g., social relationships, tags, location) and the recommendation target (e.g., to recommend contents, tags, friends, people).



Fig. 1. Recommender System architecture.

The research area of RS for MSN is still in its infancy, and few solutions have been presented in the literature, mainly aimed at optimizing content dissemination in opportunistic networks through personalized recommendations. The main difference between RS for OSN and MSN relies on the knowledge they use to select the recommendations for their final users. In OSN, RS assume to have access to a complete knowledge of all the available objects in the network (i.e., contents, tags, etc.), residing on a centralized infrastructure. Instead, in MSN, RS can rely on the local knowledge of each user, represented by the local available objects and those declared by other users in proximity through D2D communications. In this scenario, each mobile device has a different knowledge of the network, which grows up through its personal mobility and its opportunities to communicate with the others. In addition, mobile devices have limited computational capabilities and RS must be characterized by efficient response time. This is due to the limited and unpredictable duration of the opportunistic contact, during which mobile devices can exchange their local knowledge and the recommended objects. In order to evaluate and compare the performance of RS for MSN, it is necessary to reproduce the realistic behavior of mobile users in a synthetic environment, since a common evaluation framework and appropriate real datasets are not currently available. In Section 5 we present the proposed RS in this area, highlighting the advantages of using context information to improve the recommendation process and demonstrating their efficiency with respect to centralized solutions. Eventually, in Section 6 we present some concluding remarks and open research challenges in this area.

2. The recommendation task

As a general concept, RS try to identify and foresee users' interests in specific contents, based on their previous experiences. Fig. 1 depicts the high-level architecture of a typical RS. When a user interacts with the system, she provides a set of explicit or implicit feedbacks (e.g., likes, clicks, ratings) about her tastes. For example, if a user positively rates an article about a new smartphone, she may also be interested in reading news about mobile apps. Therefore, the basic idea of RS is to exploit this information to infer user interests. Based on the user's past feedbacks, RS learn a model to predict how much a user can be interested in new items. Those items are then ranked according to their predicted relevance for the user. Finally, the higher ranked items will be proactively suggested to the user.

The relationship between users and items is generally represented by a matrix $\mathbf{R}_{\mathbf{M}}$ storing a rating $r_{u, i}$ to identify how much the user u liked the item i in the past. More formally, the matrix $\mathbf{R}_{\mathbf{M}}$ is called *ratings matrix* and it can be defined as follows:

$$\mathbf{R}_{\mathbf{M}}: U \times I \to R,\tag{1}$$

where $U = \{u_1, \ldots, u_m\}$ represents the set of users, $I = \{i_1, \ldots, i_n\}$ is the set of items, and $R = \{r_1, \ldots, r_k\}$ denotes the set of possible ratings that users can use to express their grade of interest for a specific item.

The semantic value of the *ratings* highly depends on the application domain. They are often specified as a discrete set of ordered numbers (e.g., a 5-point rating scale is commonly used in ecommerce websites, like Amazon, and video-on-demand services, like Netflix), or a binary value, as in most OSN (i.e., "like" on Facebook¹ or a "retweet" on Twitter²). In the first case, the ratings are considered as *explicit feedbacks* of the users, while in the latter they are considered as *implicit feedbacks*, derived from the users' activities in OSN. In this paper, we generally refer to ratings scale including both cases.

A missing value in the ratings matrix can have two meanings: (i) the user does not want to express an opinion about a specific item, or (ii) the user does not know that item yet.

As we will discuss in details in the next sections, a great number of recommendation techniques (e.g., Collaborative Filtering) aims to predict missing ratings in $\mathbf{R}_{\mathbf{M}}$ in order to generate lists of new items to recommend to the users. However, there are also other Recommender Systems (e.g., some network-based RS) that simply provide ranked lists of items without explicitly predicting ratings, using the ratings matrix just to model the users' preferences.

Typically, there are a lot of missing values in $\mathbf{R}_{\mathbf{M}}$. In [5] the sparsity of the ratings matrix has been calculated to be generally larger than 99% in commercial systems. This is due to the *long tail* property [6], which is satisfied by different real-world settings. According to this property, just a small subset of the contents in the system is associated with a high number of rates (i.e., *popular items*) and, consequently, the vast majority of the items is rarely rated by the users. This characteristic of the ratings matrix has an important implication for RS: the system may not have enough information to generate relevant recommendations for the target user, and the presence of popular items may bias the recommendation process, likely providing trivial recommendations. The main recommendation approaches presented in Section 3 address this problem in different ways. However, before detailing them, we briefly describe the most used evaluation metrics.

2.1. Evaluation metrics

RS performance can be measured in terms of accuracy in predicting ratings or improvement of the user's experience. This choice depends on the RS target: to increase the profit of the service or to improve the user's satisfaction.

In addition, RS can be evaluated through *online* or *offline* approaches. The *online* approach mainly leverages on the continuous interaction between the users and the system; the user selects an item among those recommended, and this choice is used as input to machine learning algorithms used to adapt the RS behavior

to the user's preferences. This approach requires that a large number of users is involved in the RS evaluation, and only few research works explored this kind of evaluation until now (e.g., [7–9]). Most RS are currently evaluated using the *offline* approach by exploiting available datasets including the history of the users actions. In these cases, the datasets are generally divided into two subsets, namely the *training* (T_r) and *test* (T_e) sets. The former represents a set of examples used to fit the parameters of the proposed recommendation model, and the latter is used to test its accuracy.

2.1.1. Prediction accuracy

The prediction accuracy measures how much the predicted ratings, derived from the learning phase of the RS, differs from the actual ratings computed on the test set T_e . More formally, let $r_{iz} \in T_e$ be the explicit rating given by the user u_i for the item i_z in the test set, and \hat{r}_{jz} be the rating predicted by the RS for the user u_j related to the item i_z , the accuracy is determined by the error computed for the user-item couple $(u_j, i_z) \in T_e$ as $e_{jz} = \hat{r}_{jz} - r_{jz}$. Leveraging on the single e_{jz} errors for each pair $\langle user, item \rangle$, different accuracy metrics have been proposed to compute the overall error over the entire test set T_e . The most used are Root Mean Squared Error (RMSE) [10] and Mean Absolute Error (MAE) [10]. However, they both depend on the specific rating scale used in the reference dataset. Thus, in order to compare the performance of a RS over different datasets, we should use their normalized versions: Normalized RMSE (NRMSE) and Normalized MAE (NMAE), defined as follows:

$$RMSE = \sqrt{\frac{\sum_{(u_j, i_z) \in T_e} e_{jz}^2}{|T_e|}}; \quad NRMSE = \frac{RMSE}{r_{max} - r_{min}}, \tag{2}$$

$$MAE = \frac{\sum_{(u_j, i_z) \in T_e} |e_{jz}|}{|T_e|}; \quad NMAE = \frac{MAE}{r_{max} - r_{min}},$$
(3)

where r_{max} and r_{min} are, respectively, the maximum and minimum ratings in the dataset. Both RMSE and MAE (and their normalized versions) are meaningful accuracy measures: smaller values, better RS performance. However, as RMSE sums the squared errors, it may be highly affected by outliers. This means that a few wrong predictions could lead to a worst overall RMSE.

2.1.2. Ranking accuracy

Ranking accuracy deals with the different levels of utility of the recommended items with respect to their position in the ranked list proposed to the user. Ideally, the best ranking in the recommendation list should be assigned to the items considered as the most useful to the user. One of the most popular metrics used to evaluate the ranking accuracy is the *Discounted Cumulative Gain* (*DCG*) [11] and its normalization, *Normalized Discounted Cumulative Gain* (*NDGC*) [12], which are defined as follows:

$$DCG = \frac{1}{m} \sum_{u=1}^{m} \sum_{j \in I_u} \frac{g_{uj}}{\log_2(v_j + 1)}; \quad NDCG = \frac{DCG}{IDCG},$$
(4)

In DCG formula, *m* is the number of users in the test set T_e , I_u represents the set of items rated by the user *u*, and v_j is the position of the item *j* in the recommended list. The numerator, g_{uj} , represents the utility (i.e., the gain) of the item *j* for the user *u*, and the denominator represents a discount factor with respect to the item position in the ranking list. The utility directly derives by the ground-truth (i.e., the actual rating r_{uj} provided by the user in the test set). The Ideal value IDCG is computed with the same DCG formula by using the ground-truth ranking.

Generally, a RS suggests to the users the top-k elements of the recommendation list. An alternative way to evaluate its accuracy is to consider the trade-off between the length of the list R_L and

¹ https://www.facebook.com

² https://twitter.com



Fig. 2. Example of matrix factorization.

the number of actual relevant items for the user (i.e., the items positively rated by the user in the test set).

The relevant items contained in R_L are also identified as the *true-positives* (*tp*), while the relevant items not included in R_L are called *false-negative* (*fn*). In the same way, the proposed items not actually relevant for the target user are called *false-positive* (*fp*), while those not relevant and discarded are called *true-negative* (*tn*). Therefore, the trade-off between the length of R_L and the number of relevant items can be measured using the *Precision and Recall* measures defined as follows [13]:

$$Precision = \frac{tp}{tp + fp}; \quad Recall = \frac{tp}{tp + fn}$$
(5)

In addition, one way to create a single metric that summarizes both the aforementioned measures is the *F-measure* [14], which is the harmonic mean of equally weighted Precision and Recall:

$$F = \frac{2 \cdot Precision \cdot Recall}{Precision + Recall}$$
(6)

3. Main approaches for Recommender Systems

3.1. Collaborative Filtering

Collaborative Filtering (CF) is considered the most popular and widely implemented technique in RS [15]. The underlying assumption of CF is that people with similar preferences will rate the same objects with similar ratings [16]. Existing CF solutions can be categorized in two main classes: (i) memory-based and (ii) modelbased methods. Memory-based solutions leverage on similarities in users' behaviors and preferences to make inferences about missing values in the ratings matrix. Instead, model-based methods exploit the matrix values to learn a model, similarly to a classifier that trains a model from labeled data. The learned model is then used to predict the relevance of new items for the users.

3.1.1. Memory-based

Memory-based algorithms (also known as *Neighborhood-based*) rely on the notion of similarity among users, or items, to predict the possible interest of a user in items that she has not seen (or rated) before. In the literature, memory-based CF solutions are typically divided in two main categories: *user-based* and *item-based*. The former approach is based on the assumption that similar users typically rate the items in a similar way. Therefore, a user-based CF predicts the rating that a user u might assign to an item by aggregating the ratings that the most similar users to u have previously given to that item. Formally, the predicted rating of a user u to the item j can be formulated as follows:

$$\hat{r}_{u,j} = \frac{1}{K} \sum_{k \in N_u} Sim(u,k) \cdot r_{k,j},\tag{7}$$

where N_u is the set of the *K* users most similar to the target user *u*; Sim(u, k) represents the similarity between the users *u* and *k* for a predefined similarity measure, and $r_{k, j}$ represents the rating made by the user *k* for the item *j*.

In contrast with user-based CF, item-based CF focuses on the similarities among items. It is based on the assumption that similar items are rated in a similar way by the same user. In this case, the items recommended to the user u are ranked by aggregating the similarities between each candidate item and the items that u has rated in the past. It is possible to formulate the prediction rating for the item-based CF as follows:

$$\hat{r}_{u,j} = \frac{1}{K} \sum_{k \in N_i} Sim(j,k) \cdot r_{u,k},$$
(8)

where N_i represents the set of neighbor items of item j, and Sim(j, k) is the similarity value (for a predefined similarity measure) between items j and k.

The similarity computation between users (or items) represents a crucial step in memory-based approaches, as it can seriously reduce both the accuracy and the performance of RS [17]. In the literature, a number of similarity measures have been proposed, but the *Pearson Correlation* (PC) similarity seems to provide the most accurate results [18]. PC selects just the co-rated items and considers the differences in the mean and variance of the ratings made by two users *u* and *v*. PC similarity between two users is given by the following formula:

$$PC(u,v) = \frac{\sum_{i \in I_{u,v}} (r_{ui} - \bar{r}_u) \cdot (r_{vi} - \bar{r}_v)}{\sqrt{\sum_{i \in I_{u,v}} (r_{ui} - \bar{r}_u)^2 \cdot \sum_{i \in I_{uv}} (r_{vi} - \bar{r}_v)^2}},$$
(9)

where $I_{u,v}$ is the set of items rated by both users u and v, r_{ui} is the rating made by u for the item i, and \bar{r}_u is the mean value of the ratings made by u. On the other hand, PC similarity between two items i and j can be calculated comparing the ratings made by users that have rated both items:

$$PC(i, j) = \frac{\sum_{u \in U_{i,j}} (r_{ui} - \bar{r}_i) \cdot (r_{uj} - \bar{r}_j)}{\sqrt{\sum_{u \in U_{i,j}} (r_{ui} - \bar{r}_i)^2 \cdot \sum_{u \in U_{ij}} (r_{uj} - \bar{r}_j)^2}},$$
(10)

where $U_{i, j}$ is the set of users who rated both items *i* and *j*, and \bar{r}_i is the average rating value obtained by the item *i*. For a complete comparison of the similarity measures typically used in memory-based CF, we refer the reader to [17].

The main drawback of memory-based CF is represented by the computational cost: the computation of similarities between all pairs of users, or items, typically requires a quadratic time. However, they gained popularity due to their very simple implementation, providing an intuitive justification for the computed predictions. For instance, in item-based CF, the list of similar items can be presented to the user in order to better understand the reason of the recommendation.

3.1.2. Model-based

Although memory-based CF approaches are easy to implement and useful in effectively predicting missing ratings, model-based solutions typically show more accurate results [19]. The main characteristic of model-based CF is the use of machine learning techniques to learn models that are able to predict the missing ratings. In the last few years, different models have been proposed for model-based CF, such as Association rule-based [20,21], Bayesian

networks [22,23], Support Vector Machines [24], Neural networks [25] and, more recently, Deep learning methods [26]. However, the Matrix Factorization (MF) models [27] are considered to be the state-of-the-art in recommendation systems due to their advantages with respect to scalability and accuracy [19]. Generally, MF models exploit the typically high correlation between rows (e.g., users) and columns (e.g., items) of an incomplete ratings matrix in order to learn low-rank representations of both users and items. Moreover, these low-rank matrices (also referred as *latent factors*) can be used to approximate the full rating matrix and, then, to predict missing scores between users and items. Specifically, latent factors derived from the correlation patterns in the ratings matrix, which can be used to describe with more details both users and items profiles. Consider the toy-example depicted in Fig. 2. The ratings matrix contains the preferences of 3 users for 3 different books. The latent factors may represent the genres of the books. Therefore, the users' latent factors describe how much each user is interested in a specific genre and, in the same way, the items' latent factors represent how much a book belongs to a specific genre.

Formally, the matrix factorization method can be expressed as follows:

$$R \approx U \cdot V^T, \tag{11}$$

where *R* is the full rating matrix, and *U* and *V* are two matrices of users and items latent factors, respectively. The main goal of a MF algorithm is to learn the optimal latent factors (U^* and V^*) which better approximate *R*. The most common formulation of MF is formalized as follows [5]:

$$U^{*}, V^{*} = \operatorname*{argmin}_{U,V} \left\{ \frac{1}{2} \sum_{i=1}^{M} \sum_{j=1}^{N} I_{ij} (r_{ij} - U_{i}V_{j}^{T})^{2} + \frac{\lambda_{U}}{2} \|U\|_{F}^{2} \frac{\lambda_{V}}{2} \|V\|_{F}^{2} \right\},$$
(12)

where U_i is the latent factors of user *i* and, in the same way, V_i represents the latent factors of item *i*. I_{ij} denotes an indicator function that is equal to 1 if rating r_{ij} is not a missing value of the starting rating matrix. Finally, $\|\cdot\|_F^2$ represents the (squared) Frobenius norm of a matrix, and λ_U and λ_V are regularization parameters to set in order to prevent the model's overfitting. *U* and *V* can be seen as unknown variables, which need to be learned in order to minimize the objective function presented in Eq. 12. This is typically achieved with the well-known Gradient Descent algorithm or one of its variants (e.g., Stochastic Gradient Descent – SGD).

Compared to memory-based approaches, the predictions' accuracy is not the only advantage of model-based CF. In fact, space requirements for model-based CF are often lower than those required by memory-based algorithms. This is due to the fact that memory-based approaches should load into memory all the ratings in order to perform the recommendations, while the model learned by model-based CF is typically much smaller than the original ratings matrix. However, learning a model may require lots of training data and time. In addition, if the system is highly dynamic (e.g., new users or items are frequently added), it could be necessary to train a new model several times since the current one can easily become obsolete, thus affecting its accuracy. This could be an issue for some specific application domain. For example, in a Mobile Social Network, each node (e.g., mobile device) has a limited knowledge of the items (and users) available in the network, and it could be difficult to train an accurate model.

3.2. Content-based Recommender Systems

CF approaches can be considered inherently "social" since they exploit the correlations in the ratings patterns among users of the same system. However, CF methods typically ignore the item's attributes to compute the predictions in their recommendation process, while they can be useful to improve the recommendations' accuracy.

Content-based Recommender Systems [28,29] are specifically designed to recommend items similar to the ones that the user has preferred in the past. In CF the similarity between two items (or two users) is calculated in terms of correlation or similarity among ratings provided by other users. Instead, content-based approaches consider only the ratings provided by the target user and the features of the rated items [19]. Content-based RS are typically implemented using traditional classification and clustering algorithms, such as Support Vector Machines (SVM) or Nearest Neighbors methods (NN) [29]. In SVM the features of the items and the ratings made by the target user are combined together to form a dataset of *< features, rating >* instances. Based on this dataset, a specific classifier (or regressor) for the target user is trained in order to predict the ratings of new items, never seen before by the user. In NN, in order to classify a new, unrated item, the algorithm compares it with all the stored items, derived from the training data, by using a similarity function (typically the Euclidean distance or the Cosine similarity).

According to Aggarwal [19] and Tang et al. [16], content-based RS have the following drawbacks:

- 1. *Items' features*. The accuracy of RS relies on the set of features that describe the items. The identification of the most relevant features is not trivial, and highly depends on the specific application.
- 2. Over-specialization. Since the content-based methods rely just on the characteristics of the items already rated by the target user, it typically suffers from low diversity and novelty of the recommended items. Thus, a content-based RS typically tends to provide obvious rec-

ommendations, which could bother the user in the long term.

3. *Training set size*. In order to allow a content-based RS to learn the user's preferences, the user has to rate a sufficient number of items [16]. Otherwise, the RS does not has enough information to learn an accurate model and it fails to recommend items for a user with few or no ratings.

Despite the aforementioned drawbacks, content-based RS are useful to alleviate some critical problems of CF. For example, when a new item is added to the ratings matrix, CF methods are not able to perform recommendations about this item because the system has not yet collected sufficient ratings about it. In the literature, this problem is called *cold-start*, and a content-based approach can well complement a CF system because of its ability to exploit the characteristics of the items in the recommendation process [19,30]. Mixing two or more approaches is often referred as *Hybrid RS* [31] in which the goal is to combine the strengths of different methods to create a more robust RS.

3.3. Graph-based

Social Networks can be naturally modeled as graphs, in which nodes represent users and items, and edges model the different relationships among user-user or user-item pairs (e.g., friendships, follow, likes, share). These relationships can also be used in RS to identify similar users and/or items. In fact, several graph-based RS have been recently proposed in the literature [32]. Their first objective is to perform a ranking of the nodes to identify the most interesting items for a user or a neighborhood (e.g., the most similar users/items) to be used in a traditional CF method.

PageRank [33] is clearly the most famous method to produce a ranking of the nodes in a graph. It is based on the idea that a node is important if it is linked to other important nodes. This simple

and general idea can be applied to many different situations, including ranking of websites in a search engine, or finding relevant items (or people) in a RS. PageRank exploits a random-walk model. Specifically, it performs an exploration of the given graph by visiting nodes following randomly selected links among them. The long-term frequency of visits to a node represents its final ranking value, which is clearly influenced by the number of incoming links of the node. It is also referred to as the steady-state probability. However, directed graphs can have some nodes (or group of them) without out-going links; this part of the network would act as a trap for the random-walk process. In order to overcome this problem, PageRank uses the so called *restart* mechanism. According to this strategy, at each transition, the random walk may either jump to an arbitrary node of the network with a probability α , or follow one of the out-going links connected to the current node with probability $(1 - \alpha)$. Formally, the steady-state probability at a node *i* is defined as follows:

$$\pi(i) = \frac{\alpha}{n} + (1 - \alpha) \cdot \sum_{j \in ln(i)} \pi(j) \cdot p_{ji},$$
(13)

where *n* is the total number of network's nodes, In(i) is the set of nodes that have out-going links directed to *i*, and p_{ji} is the transition probability from node *j* to node *i*.

Although PageRank is very effective in finding popular nodes in the graph, it does not take into account the user's preferences, which means that it is not able to provide personalized recommendations. To this aim, Haveliwala [34] proposes personalized PageRank, a biased version of PageRank, specifically designed to find popular nodes which are also similar to specific nodes in the graph. In this case, the key idea is to use the restart mechanism to bias the random walk process towards specific nodes that represent the user's interests. If we consider a specific OSN represented as a graph, nodes represent the items available in a social media, and the links represent the relations between users and items. The restart nodes represent the items consumed by the target user in the past. Then, exploiting the biased random-walk, personalized PageRank provides a way to calculate similarity scores of nodes based on their structural similarity to the restart nodes. Finally, the items with higher similarity values can be recommended to the target user.

Even if personalized PageRank, and other similar random walkbased algorithms (e.g., *FolkRank* [35], *SocialRank* [36], and *Item-Rank* [37]), can be directly employed to build Social Recommender Systems (SRS) [36,38], they are mostly used to select user/item's neighbors in memory-based CF solutions [39,40]. In fact, rather than using standard similarity measures (e.g., Pearson Correlation Coefficient on user's ratings), it is possible to use such structural measures to identify the most relevant users (or items) for the recommendation purpose. Since they only leverage on the structural transitivity of the network's edges, instead of the users' ratings, they typically result very effective to alleviate the problems related to the sparsity of the ratings matrix. However, random walk-based approaches typically require multiple iterations before the transition probabilities converge to a steady state.

On the contrary, other solutions require only few steps to generate effective recommendations [41]. *Diffusion-based* methods [32] represent the most used approach for recommendation in graph-based RS. They have been initially proposed by Zhou et al. [41,42] with their algorithms *ProbS* (Probability Spreading) [42] and *HeatS* (Heat Spreading) [41]. These algorithms exploit a bipartite user-item graph where users are directly connected to the items collected in the past.

In ProbS, a resource is initially assigned to each item connected to the target user, initialized to a unitary value. This resource is then spread uniformly from the collected items to the users connected with them (*first diffusion step*) and then, in a second step back, to the items connected with those users (*second diffusion step*). The final value of this resource, associated with each item, is then interpreted as the recommendation score of that item for the target user. The higher the score obtained by an item, the greater could be the interest in it by the target user.

HeatS [41] is similar to ProbS, but it is based on opposite rules: each time the resource (or a portion of it) is redistributed, it is divided by the number of edges connected to the node towards which it is heading to. Fig. 3 depicts the diffusion steps of the two algorithms, highlighting the differences in the two recommendations.

Unfortunately, both of them are actually biased by the presence of extremely popular or non-popular items (i.e., nodes with very high or low degree, respectively), not taking into account the characteristics of the users' interests. To overcome these limitations, a ProbS+HeatS hybrid approach, simply called *Hybrid*, has been proposed in [41]. Hybrid calculates a linear combination of the results of ProbS and HeatS with a parameter λ governing the relative importance of one of the two original algorithms. By tuning λ appropriately, Hybrid is able to obtain simultaneous gains in both accuracy and diversity of recommendations with respect to the previous solutions.

The diffusion-based approach has been generalized in multiple ways. Some authors proposed to compensate the HeatS's preference for unpopular items modifying its spreading process [43] and allocating higher weight to links connecting nodes with high degree [44]. On the other hand, in order to compensate the ProbS's bias towards very popular items, Lü and Liu [45] proposed to highlight less popular items in the second diffusion step by making their score inversely proportional to the degree of the nodes.

Furthermore, Zhou et al. [46] investigated the effect of biasing the initial resource allocation in ProbS. Results indicate that decreasing the initial resource allocated on popular objects can further improve the algorithmic accuracy. Similarly, Liu and Zhou used heterogeneous initial resources in Hybrid, further improving both the accuracy and diversity of the recommendations [47]. However, in addition to the ProbS–HeatS hybridization parameter λ , they introduce a second parameter to tune, which governs the heterogeneity of the initial resources.

For a complete comparison and evaluation of the diffusionbased methods, we refer the reader to [32].

3.4. Context-Aware Recommender Systems

RS we have presented focus on recommending the most relevant items to individual users taking into account just two entities: users and items they liked in the past. However, they do not consider additional information that can be useful to improve the recommendations' quality and personalization [5]. To this aim, *Context-Aware Recommender Systems* (*CARS*) [48] have been recently proposed in the literature. CARS methods perform their recommendations considering also the information characterizing the specific situation in which the target user is involved. This additional information is referred to as the *context*. The notion of context has been studied in multiple disciplines and several heterogeneous definitions exist [48]. In CARS, context is defined as the additional information relevant to improve recommendations, such as the time of the day, the locations of the target user, and her social relationships.

As we described in Section 2, the objective of a traditional RS is to approximate the rating's function $f_R: U \times I \rightarrow R$ in order to predict the possible interest of a user for a specific item. CARS generalize this approach by using a *multidimensional approach* [49] in which the rating function can be seen as a mapping from a *n*dimensional tensor (i.e., a multidimensional matrix) to the set of



Fig. 3. Execution of ProbS and HeatS on a bipartite user-item graph.



Fig. 4. Example of a multidimensional cube for the *User* × *Item* × *Location* recommendation space.

ratings R:

$$f_R: D_1 \times \dots \times D_n \to R, \tag{14}$$

where two dimensions, typically, are the sets of users and items (i.e., *U* and *I*), and the other dimensions represent the context variables. Grafically, CARS's rating function can be represented as a multidimensional cube. Fig. 4 shows an example with a 3-dimensional cube that stores the ratings for the recommendation space User × Item × Location, where $f_R(u_1, i_4, home) = 5$ means that the user u_1 rated with a score of 5 the item i_4 when she was at home.

The paradigms used to incorporate context information in RS are divided in three main categories [50]: (i) *Contextual pre-filtering*, (ii) *Contextual post-filtering*, and (iii) *Contextual modeling*. In the following we describe in details each approach.

Contextual pre-filtering. In contextual pre-filtering the idea is to reduce the multidimensional cube to a 2-dimensional matrix in order to apply traditional algorithms. As depicted in Fig. 5, the current context c of the target user is used as a filter to "slice" the rating cube and extract the relevant user–item matrix, which can be used with traditional CF algorithms.

However, using the exact context c sometimes can be a problem. If there are just few ratings associated with c, this can lead to a low accuracy in the recommendation because the system has not enough data points about the past preferences of the target user. Therefore, the notion of generalized pre-filtering [49] has been introduced in the literature to cope with this problem. Instead of using the specific context c, this approach suggests to use a broader context c', which includes c and other similar context variables. Let us consider the context of watching a movie on Saturday (i.e., c = "Saturday"). In order to base the recommendation on a wider set of ratings, the generalized pre-filtering approach may use the more general context c' = "Weekend", which includes the two context variables c = "Saturday" and $c_1 =$ "Sunday". Then, the system aggregates the selected ratings using some aggregation function (e.g., the average) to reduce the recommendation space to the 2dimensional user-item matrix and, finally, perform the recommendation.

It is clear that the context generalization becomes a crucial aspect of the recommendation process. One possible strategy is to use a process similar to the cross-validation: in a learning phase, the predictive performance of the system is empirically evaluated using different contextual filters. Then, the filter that allows the best performance is automatically selected [50]. In addition, multiple context generalization can be used for the same specific context. For this reason, in the literature several works propose to use a combination of different filters and prove that this solution can provide a significant performance improvement compared to the use of a single filter [50,51]. In this case, multiple pre-filters are used to generate several different 2-dimensional ratings matrices. Eventually, those user-item matrices are merged in one single matrix in order to perform the recommendation process. Codina et al. recently proposed in [52] the use of semantic similarities between contextual situations in order to select the most relevant context variables for a specific situation. They showed that their semantic pre-filtering model is able to obtain better results than others context-aware approaches.

Contextual post-filtering. In this case the filtering step is applied to the output obtained after applying traditional RS to the data set. Therefore, the recommendation process of a post-filtering method can be summarized in two steps:



Fig. 5. Process of a pre-filtering method to generate contextual recommendations.



Fig. 6. Process of a post-filtering method to generate contextual recommendations.

- 1. In a first stage, the context information are completely ignored, and a list of recommended items are generated using traditional RS on an aggregated 2-dimensional user-item matrix.
- 2. Given a specific context *c*, recommendations not relevant for *c* are filtered out, and the ranking of recommendations in the list is adjusted according to *c*.

Fig. 6 depicts the recommendation process of the post-filtering approach. Here, aggr(U, I) refers to the aggregation function (e.g., the average) used to create the 2-dimensional user–item matrix, and *filter*(*c*) represents the filtering (or adjustment) step of the recommendation's list, based on the current context *c* of the target user.

The selection of the best contextual filter represents a crucial point in the overall recommendation process. All the techniques related to the context generalization, previously discussed, can be applied also in post-filtering methods. The main difference between the two approaches lies in the data set on which the contextual filter is applied: the ratings contained in the multidimensional tensor (i.e., the input) for the pre-filtering approach, and the recommendation list (i.e., the output) generated by a traditional RS for post-filtering methods.

Both approaches share the same advantages: their implementation is relatively simple, and both of them allow the use of the traditional recommendation techniques described in the previous sections. In addition, comparing their performance, none of them seems to clearly outperform the other. Panniello et al. [53] evaluated different solutions based on both approaches by using several real world datasets. The authors did not find any solution that dominates over the others for all the datasets, indicating that the choice of the best approach really depends on the specific application.

Contextual modeling. Contextual modeling approaches consider context information directly in the recommendation function in order to predict a user rating for a specific item [50]. Differently from pre- and post-filtering techniques, contextual modeling methods use predictive models or heuristics to create multidimensional recommendation functions.

The simplest solution adapts neighborhood-based CF methods to perform context sensitive recommendations [49]. In this case, the traditional user–user (or item–item) similarity metric is replaced by an *n*-dimensional distance metric, which includes the contextual information. For instance, let's consider two points in a 3-dimensional user–item–context cube, respectively A = (u, i, c)and B = (u', i', c'). The distance between them, Dist(A, B), can be calculated as the sum of the weighted distance between each dimension or, alternatively, using the Euclidean metric [49]. Therefore, by extending the traditional neighborhood rating prediction (i.e., Eq. (7)), the specific rating for user *u*, item *i* in the context *c*, can be expressed as follows:

$$r_{u,i,c} = k \sum_{(u',i',c') \neq (u,i,c)} \frac{1}{Dist((u',i',c'),(u,i,c))} \cdot r_{u',i',c'},$$
(15)

where *k* is a normalization factor, and the inverse of the multidimensional distance between points has been used for weighting the rating $r_{u',i',c'}$ (i.e., the more the two points are close, and the higher is the weight).

More complex approaches extend model-based solutions. In particular, tensor factorization methods [54,55] can be considered as a generalization of matrix factorization for latent factor models (see Section 3.1.2). As depicted in Fig. 7, a 3-dimensional tensor factorization aims at learning three latent factors matrices \hat{U} , \hat{l} , \hat{C} (i.e., respectively, the user, item, and context latent factors) and a core tensor \hat{T} to build a prediction model. However, the computational complexity of this kind of models increases exponentially with the number of the tensor's dimensions [56]. To solve this issue, some simplified forms of these factorization methods have been recently proposed. For instance, the *Pairwise Interaction Tensor Factorization (PITF)* [56] and *Factorization Machines* [55,57] are able to predict missing ratings exploiting the pairwise interactions between the different dimensions of the tensor, maintaining at the same time a low computational complexity.

To conclude the description of main RS methods, we present in Fig. 8 a summary of their classification and the references of the main representative solutions for each method.



Fig. 8. Classification of Recommender Systems methods.

4. Recommender systems for Online Social Networks

With the advent of OSN, RS have been further enhanced by exploiting additional information characterizing both users and items, becoming thus SRS [16,58]. Specifically, OSN can provide information able to identify virtual and physical social relationships among users, common interests and habits, in addition to personal preferences. RS can exploit this information depending on the application domain in which they are implemented. Specifically, SRS widened their recommendation targets, from generic item recommendations up to people/friend recommendations, locations, tags and others. In this section, we provide a description

of SRS grouped by the type of recommendations they provide, and the context information used to optimize their process, as shown in Fig. 9. Specifically, we identify three types of recommendations: (i) *social-aware recommendations*, in case SRS recommend people or friends, and include the notion of trust and social relationships in generic items recommendation; (ii) *tag-based recommendations*, when SRS exploit a specific characterization of items based on *tags* to recommend other items, tags and/or people, e.g., if they share similar interests derived from tags; (iii) *location-aware recommendations*, when SRS exploit location-related information to recommend items, point-of-interest (POIs), trajectories, and people.



Fig. 9. Classification of Social Recommender Systems.



Fig. 10. Graphical representation of (a) a basic friendship network, and (b) a trust-network with inferred trust links (dashed lines).

4.1. Social-aware recommendations

OSN provide several information that characterizes the social context of a user and her generated items. OSN analysis mainly refers to the topological structure of the social network, in terms of (virtual) social relationships among users [59,60]. In addition, OSN can reflect the presence of the homophily principle [61] also in the cyber-world, as the tendency of individuals to associate and connect with similar others, as demonstrated by recent studies [60].

The social structure of a OSN can be represented as a graph G = (U, S), where U is the set of users and S is the set of social links among them, which can be: (i) undirected, if they model friendship relations, such as in Facebook, or (ii) directed, if they model trust or follow relations, such as in Epinions³ or Twitter. Social links in OSN represent the explicit declaration of a user to be virtually in contact with others, but they do not reveal anything about the effective nature of the relationships or its strength and value for the user. However, analyzing OSN contents and users' virtual interactions (e.g., explicit users' feedbacks, ratings [62], etc.) we can derive additional information that can further characterize those links, like for example the notion of *trust* among users, which can be represented through links' weights (Fig. 10(b)).

SRS exploit this structure to provide two types of social-aware recommendations: (i) friends or other people, and (ii) OSN items.

In the first case, SRS suggest to the users new social relationships based on the prediction of common interests, while in the second case they exploit social/trust relationships in order to suggest interesting items to the users. In the following we review the main solutions proposed in the literature, highlighting different approaches, and how they take into account different aspects of the social context.

4.1.1. Friends and people recommendations

This recommendation task differs from the generic items recommendation since it must take into account specific aspects of the social network of each user, like trust, reputation, privacy and personal attraction [58,63]. From a technical point of view, this task is associated with the problem of link prediction in the social network, which tries to infer new possible relationships or interactions between pairs of entities, based on their properties and the existing links [64].

In the literature, several techniques have been proposed for link prediction [65]. The most commonly used can be divided into two main categories: unsupervised and supervised methods. The first category includes all those techniques that exploit the structure of the network to calculate similarity values among the users (nodes of the graph), and then recommend the most similar ones. For instance, Jaccard coefficient [65] or the Adamic–Adar measure [65] can be used to calculate the similarity between two nodes based on their common neighbors. Alternatively, random walkbased approaches such as Katz measure [65], or PageRank [33] and

³ http://www.epinions.com

SimRank [66] algorithms, can represent a different way to recommend new links based on the connectivity between pairs of nodes.

One of the first unsupervised RS for people recommendation has been proposed by Guy et al. for an IBM enterprise social network called SONAR [67-69]. It defines a "relationship score" based on a combination of various social context information derived from the enterprise intranet applications (e.g., organizational charts, papers database, patents database, a social networking site, a friending system and others). For each application, they define a user-user graph in which the links' weight is computed with a specific metric depending on the application (e.g., manageremployee relationship, paper/patent co-authorship, users belonging to the same group in the friend network). Then, a linear combination of the weights of each graph is computed for each pair of users, and it is used as general relationship score. The authors compared SONAR relationship score with that provided by other approaches: (i) Friend-of-friends (FOF) metric, applied to the internal friending system of the company; (ii) a Content-Matching approach, defining the users' similarity based on a set of keywords extracted only from the contents generated on the social network (i.e., posts, picture descriptions and tags); (iii) Content-plus-Link, by doubling the similarity value derived from the Content-Matching in case the pair of users are also explicitly connected in the social network application. They evaluated SONAR with real experiments involving up to 3000 users in different testing scenarios and demonstrated the effectiveness of the proposed solution with respect to the other approaches.

Other solutions have been proposed by exploiting the concept of *ego-network*, the representation of a social network centered on the individual (i.e., ego) and her social links (i.e., alters) divided in a series of "circles". The inner-most circle includes alters with a very strong relationship with the ego. Each subsequent circle (in hierarchy) includes all the relationships of the previous circles along with an additional set of social links characterized by a weaker level of intimacy. The last set, included in the outermost circle only, contains simple acquaintances, with a relatively weak relationship with the ego. This representation has been defined in antropology for classical (offline) social networks [70,71], and it has been recently demonstrated that the same representation can also be adopted in OSN [72,73].

Epasto et al. [74] proposed a RS for friends recommendation based on the features derived from the co-occurrences of two nodes in different ego-networks. In contrast to random walk-based measures, these features can be computed much more efficiently on very large graphs, by just analyzing selected neighbors of each node. The authors evaluated the system on an anonymized snapshot of Google+ social network, and compared the proposed solution with two other unsupervised methods (i.e., common friends and Adamic–Adar) in terms of precision and recall. A 2-week test, during which the authors observed the social network evolution based on the different recommendations, demonstrated the effectiveness of their solution.

In general, the effectiveness of the proposed unsupervised methods highly depends on the social network structure used for their evaluation, and they are not easy to compare [75]. On the other hand, supervised methods try to cope with this issue, by treating the link prediction in a social network as a classification problem and by defining the characteristics of the network in a data-driven way, i.e., by creating a vector of features for each pair of nodes, in which the set of features can include different heuristics and metrics. To this aim, for each pair of nodes, a multidimensional vector is extracted and associated with a positive or negative label, based on the presence or absence of a link between the nodes. Finally, a classifier is trained using the set of vectors as training data. In this way, the classifier automatically learns the relevance of the features, and the learned model can be used to predict the creation of new links (i.e., the label) between any pair of nodes.

The following solutions are representative of this class of link prediction approaches. Fire et al. [76] proposed a set of features related to different characteristics of the graph (e.g., nodes, edges, subgraphs, and multi-hop paths) in order to learn a classifier that identifies the missing links. They used different datasets (e.g., Facebook, Flickr, YouTube, Academia, etc.) and by using a common framework they developed different classifiers for each dataset (e.g., decision tree C4.5, kNN, Naive Bayes, SVM, neural networks) in order to evaluate both single solutions and the combinations of some of them. Finally, they compared the prediction accuracy by using all the features or some limited sets.

Backstrom and Leskovec [77] proposed Supervised Random Walks, a supervised algorithm that combines the network structure with attributes about both nodes and edges in a unified link prediction algorithm. Given a source node s, and training examples about which nodes s will connect to in the future, the algorithm learns a function that assigns a strength (i.e., random walk transition probability) to each edge, providing higher scores to the possible new links of s. This solution has been compared to both unsupervised methods (e.g., random walk with restart, Adamic-Adar, Common Friends) and other two supervised approaches (i.e., decision-tree e logistic regression) on heterogeneous datasets (i.e., arXiv and Facebook). All the supervised methods outperformed the unsupervised ones in terms of precision and the ranking metric AUC [108], and Supervised Random Walk performed better than the others of the same class. We can thus observe that supervised methods generally performs better than unsupervised ones in friends and people recommendations, thanks to their ability to automatically learn the hidden relationships between social links and their features.

4.1.2. Items recommendations

Social context positively impacts also on generic items' recommendation (such as books, movies, or news), and several works have been presented in the literature exploiting most of the standard recommendation techniques. The first category we consider is related to the neighborhood CF approach. These solutions use various strategies to select relevant users for the recommendations purpose. For instance, authors in [78,79] considered only the connected friends of the target user and their items as the possible objects to recommend. Instead, in [79] an evaluation between different strategies has been proposed. Specifically, the authors considered (i) a Friends Average approach, which simply averages the ratings of the direct friends, (ii) a Weighted Friends, in which cosine similarity among friends is used to define links' weight, and (iii) a traditional CF method without social context. All the socialaware recommendation tasks achieve better results than traditional CF in terms of prediction accuracy.

Other memory-based solutions exploit trust relations among users to obtain the set of relevant neighbors [80–82]. Generally, these solutions exploit different graphs' exploration techniques in order to propagate trust within the user-to-user network. In fact, in this case, the network is represents as a direct graph in which nodes represent the users and edges connect two users if they are friends or at least one of them declares a trust towards the other greater than 0. As a first step, the trust values among directly connected users are typically calculated leveraging on a user-to-user similarity measure (e.g., the cosine similarity calculated on the ratings of the two users) [16]. Then, these values of trust are propagated through the graph's edges in order to calculate the transitivity trust among nodes that are not directly connected to each other. For instance, *TidalTrust* [80] aims at estimating the trust values of a node (i.e., source) towards another one not directly connected (i.e., sink). In a first phase, it finds a path from the source to the sink while rating the nodes on the path, and then it aggregates the trust value backwards towards the source. As a result, TidalTrust assumes that: (i) shorter propagation paths produce more accurate trust estimates, and (ii) paths with higher trust values create better results in selecting neighbor nodes to be used for items recommendations. Therefore, TidalTrust selects as relevant neighbors only those nodes that are on the shortest path between the source user and the sink. *MoleTrust* [81] is another example, which uses a random walk to propagate trust values on the graph, and it defines a threshold on the length of the path between a source and a sink. Therefore, it limits the number of possible paths between a pair of users.

In contrast with memory-based approaches, model-based Recommender Systems learn models that predict the user's interest for new items. As it happens for traditional RS, the matrix factorization (MF) represents the most used model-based technique in SRS. The common idea of Social-aware MF solutions is that social connections mainly influence the users' preferences (i.e., the ratings). The preferences of two friends should be much more similar than those of two strangers, and for this reason these solutions typically associate a weight to the social relations, which indicates the strength of the social tie among the related users and it is usually calculated based on the ratings similarity.

Each of the MF solutions proposed in the literature differs from the others by the method used to include social and/or trust information into the factorization process. For example, SoRec [83], instead of factorizing only the user-item matrix, as in traditional MF methods, it extends the factorization objective function performing a co-factorization both of the user-item matrix and the user-user social matrix. Instead, other solutions (e.g., SocialMF [84] and Social *Regularization* [85]) focus on the regularization term of the factorization objective function (see Section 3.1.2), introducing new constraints related to the social factors. Specifically SocialMF [84], by assuming that the behavior of a user is affected by the behavior of her direct neighbors, introduces a new regularization term that forces the preference of a user to be closer to the average preference of the users in her social neighborhood. On the other hand, Social Regularization [85] considers the fact that the target user's neighbors may have different tastes. For this reason it uses a pairwise regularization of the users' latent factors where the preferences' similarity of two connected users depends on their ratings' similarity. In order to evaluate Social Regularization, authors compared the performance of their solution with other three MF approaches: two traditional user-item MF (i.e., non-negative matrix factorization [86], probabilistic matrix factorization [87]), and Social Trust Ensemble (STE) [88], a social-aware MF solution where the missing ratings are predicted as a linear combination of ratings from the target user and her friends. By using the datasets extracted from two real-world OSNs (i.e., Douban⁴ and Epinions), Social Regularization demonstrated to outperform the other approaches in terms of both MAE and RMSE metrics. In addition, in the same work [85], authors propose the use of two different similarity metrics for the pairwise regularization: the Pearson Correlation Coefficient and Vector Space Similarity [89]. However, the evaluation results show no significant differences between the two measures, but they both perform better than random similarity values, showing the importance of similarity function enhanced with social context information.

Recent works (i.e., [90,91]) use different regularization terms. In [90], Sun et al. propose *RSboSN* (Recommender System based on Social Networks), which leverages on a clustering algorithm to identify groups of friends. RSboSN outperforms SoRec in terms of precision and recall, proving the effectiveness of their clusteringbased regularization approach. Moreover, Reafee et al. [91] propose the *EISR* (Explicit and Implicit Social Relation Probabilistic Matrix Factorization) algorithm, aimed at including both explicit and implicit social relations as regularization parameters. Explicit social relations are the direct friendships ties among users, while the implicit ones refer to not directly connected users and they are inferred using link prediction. Compared with Social Regularization [85] and two diffusion-based algorithms (i.e., HeatS [41] and Hybrid [41], described in Section 3.3), EISR shows a clear improvement in terms of prediction error (i.e., MAE and RMSE), demonstrating the effectiveness of taking into account also the implicit relationships in the social recommendations.

4.2. Tag-based recommendations

Over the past few years the use of special user-defined keywords, called *tags*, to categorize or describe web and OSN contents, gained a lot of popularity. This user-driven phenomenon is known in the literature as *folksonomoy* [92] and it is a wellstudied topic in both information retrieval and Recommender Systems fields [93–95].

An important aspect of folksonomies is that, differently from ontologies, no relationship between the terms is required a priori. On the contrary, relationships are automatically built exploiting the tags created by the users and explicitly assigned to contents (e.g., hashtags assigned to tweets). Since tags are generally defined as keywords that can reflect a semantic meaning of the associated content, they represent an important feature in the content characterization. In addition, folksonomies dynamically adapt to changes in the users' vocabulary, further personalizing the user's interests, and they are independent from the type of item they describe. For instance, the features used to represent the characteristics of a movie are typically very different from the features used to describe music files or books. Instead, tags can be seen as generic features that can be used to create multi-domain Recommender Systems [96].

However, folksonomies have some drawbacks due to the user language. Synonyms, homonyms, polysemies, and different users' tagging behavior make the recommendation process difficult to perform in some cases, and undermine the use of simple tag matching (e.g., group items characterized by the same tag). To overcome these limitations, a novel family of Recommender Systems, called Tag-based Recommender Systems (TBRSs) [94,95], have been proposed in the literature.

Formally, the folksonomy can be modeled in the following two ways. First, it is possible to use three distinct sets of elements: the set of users $U = \{u_1, ..., u_n\}$, the set of items $I = \{i_1, ..., u_m\}$, and the set of tags $T = \{t_1, ..., t_k\}$. These elements can represents the set of nodes in a tripartite graph, as shown in Fig. 11(a). Each binary relation between them can be described using adjacency matrices, A^{UI} , A^{IT} , and A^{UT} for user-item, item-tag and user-tag relations respectively. If the user u_l has collected the item i_s , we set $a_{ls}^{UI} = 1$, otherwise $a_{ls}^{UI} = 0$. Similarly, we set $a_{sq}^{IT} = 1$ if the item i_s has been tagged with t_q , and $a_{sq}^{IT} = 0$ otherwise. Furthermore, we set $a_{lq}^{UI} = 1$ if u_l owns items tagged with t_q , and $a_{lq}^{UT} = 0$ otherwise. Alternatively, the folksonomy can be modeled as a third-order tensor taking into account only complete ternaries Y = (u, i, t), as shown in Fig. 11(b). Specifically, the tensor's component is set Y = 1 if the ternary exists in the folksonomy, and Y = 0 otherwise.

A variety of different approaches have been proposed in the literature for tag-based RS. Works presented in [97,98] investigated the tag co-occurrence frequency to support the user tagging process. Specifically, given a specific item (e.g., a photo in Flickr), the system recommends to its users the most relevant tags to

⁴ https://www.douban.com



Fig. 11. Representation of a folksonomy with (a) a graph, and (b) a tensor.

associate with it. The underlying idea is that, if the co-occurrence frequency of two tags is high, they could be closely related. On the contrary, if two tags are not related at all, their co-occurrence frequency should be very low.

Several works extend CF towards tag-based RS. To this aim, the user-item-tag ternary is typically reduced to two-dimensional matrices. Firan et al. [99] leverage on the user-tags projection matrix to compute a ranked list of tags, then a list of recommended items is extracted according to those tags. Instead, Guy et al. [100] select neighbors users by using a similarity measure based on common tags and items, while [101] exploits the co-occurrence of tags.

In order to overcome the issues related to the folksonomies (e.g., synonyms and polysemies), and to improve the performance of CF, Shepitsen et al. [102] and Gemmell et al. [103] use hierarchical tags clustering techniques, while [104] weights each tag based on the relationships between users, items, and tags. Recently, a hybrid framework for both item and tag recommendations has been proposed in [96]. Authors use text analysis techniques (e.g., bigrams) to solve the ambiguity problem of folksonomies and perform tag recommendations. In addition, they exploit temporal information and trust relationships among the users to provide item recommendations using a CF approach.

By modeling the folksonomy as a third-order tensor, some researchers proposed Recommender Systems based on tensor factorization techniques (see Section 3.4). Specifically, in [105] the recommendation is performed by the product of three low rank matrices (i.e., the latent components of the users, items, and tags) and a low rank core tensor produced by the Higher-Order Singular Value Decomposition (HOSVD) [106]. Rendle et al. [107] proposed a different learning approach for tensor factorization models called Ranking with Tensor Factorization (RTF), which optimizes the model parameters for the ranking statistic AUC [108].

Another natural way to represent the folksonomies is to use hypergraphs. According to this representation, nodes of the graph represent the folksonomy entities (i.e., users, items, and tags), and the relationships between them are modeled with edges among nodes. Consequently, a number of researchers proposed different graph-based solutions for tag-based RS. These approaches exploit the structure of the folksonomy hypergraph in order to provide tags or items recommendations to the user. Hotho et al. [93] proposed FolkRank, a graph-based ranking algorithm inspired by PageRank [109]. Here, the underlying idea is that an item tagged with "important" tags by "important" users becomes "important" by itself. Essentially, FolkRank identifies the most "important" tags (i.e., the most visited tags during the random walk) for a specific user by performing a biased random walk with restart on the folksonomy graph. Other solutions refer to the diffusion-based approach described in Section 3.3. Zhang et al. [110] first proposed a diffusion-based Recommender System based on the tripartite user-item-tag graph. Specifically, they apply the diffusion process proposed in [42] on both the user-item and item-tag bipartite networks and then, they combine the two results to provide the recommendations to the user. Their results show that the integration of tags can enhance the accuracy of the recommendations with respect to previous diffusion-based algorithms. However, due to their diffusion process, the solution proposed in [110] and other similar works tend to suggest the most popular items in the folksonomy network, penalizing the personalization of the recommendations [41].

To this aim, we recently proposed PLIERS (Popularity-based ItEm Recommender System), a tag-based RS designed both for bipartite [111] and tripartite graphs [112]. In order to provide even more personalized items to the users, PLIERS assumes that the popularity of an item/tag (i.e., the degree of the node) can be related to its semantic. Therefore, a very popular item or tag can semantically relate to a more "generic" topic compared to a less popular item/tag that, instead, describes a more "specific" topic. For instance, any content related to the football club Millwall can be tagged with both tags "Millwall" and "Football", but the opposite is not always true: all content concerning football will not always be tagged with "Millwall". According to this assumption, we can therefore say that the tag "Football" refers to a more generic topic than that referred by the tag "Millwall". Users interested in the Millwall football club, but not connected to items tagged with "Football", are clearly not interested in all the items tagged with the latter tag, as these could contain information about other football clubs. In this ways, PLIERS solves the dilemma of the choice between popular or non-popular items in the network in a more natural way than the other diffusion-based algorithms. PLIERS does not require any tuning parameter, and it ensures that the popularity of recommended items is always comparable with the popularity of items already adopted by the users. Compared with other diffusion-based solutions, PLIERS shows comparable performance in terms of precision and recall, while providing better novelty in the recommendations.

A very recent and promising approach to build tag-based Recommender Systems is represented by the use of deep learning to discover latent features from the tag space. The system proposed in [113] uses a sparse autoencoder [114] to extract a set of dense latent features from the user-tag matrix. Based on the extracted features, users' profiles are updated and the traditional CF is used to recommend items or tags. On the other hand, Xu et al. [115] proposes a pure deep learning model to provide item recommendation to the user. Here, two neural networks map the tag-based user and item profiles to an abstract deep feature space. Then, the relevance of the specific item for the user is calculated based on the similarity between the deep representation of the two profiles. Authors compared their solution with three different approaches: a hierarchical clustering model [102], CF and the autoencoder-based model proposed in [113]. With respect to these solutions, experimental results show that the proposed approach significantly improves the effectiveness of the recommendations in terms of both precision and recall.

4.3. Location-aware recommendations

The recent advances in location-acquisition techniques, and the wide spread use of GPS-enabled smartphones, led to the creation of the so-called Location-Based Social Networks (LBSN), such as Foursquare⁵ or Yelp⁶. These systems allow users to tag, rate, and

⁵ https://foursquare.com/

⁶ https://www.yelp.com



Fig. 12. Data in location-based social networks.

describe the locations they visit, in order to aid the discovery of unknown points of interests (POIs), like restaurants or shops, or to enable new social relations among users [116].

According to Bao et al. [117], the location is one of the most important information to describe the user's context. The addition of spatial information in social networking systems is able to bridge, at least partially, the gap between cyber and physical social relationships among users. Fig. 12 depicts the main information modeled by a typical LBSN. In the upper layer, users and their social relationships are represented. These social relationships can be explicitly defined by the users (e.g., friends or followers) or they can be inferred from their activities on the OSN. The middle layer depicts the social media content generated by the users. The nature of these content can be very different; for instance, a user can share a check-in at a public place, a geo-tagged photos, a review of a restaurant, or the workout tracked by a fitness application. As depicted in the lower layer, each content should be associated with one or more geographical locations.

Given the multi-dimensional nature of the LBSNs, SRS applied to these specific OSN must be able to handle a multitude of heterogeneous data. They can be divided into four different categories based on their recommendation task: (i) location, (ii) user, (iii) activity, and (iv) social-media. The first category mainly focuses on suggesting new locations to the user, based on her preferences. The objective of the recommendations can be both single POIs or a sequence of locations, such as the best route to take or the sequence of attractions to visit during a trip. To recommend a single place, both content-based and CF techniques have been exploited. In the former case, the solutions proposed in the literature, e.g., [118–120], typically suggest new locations by matching the user's profile with the places' meta-data (e.g., text description, categories, or tags). These solutions inherit the main advantages and drawbacks of the content-based approach: they are effective in mitigating the cold-start problem but, on the other hand, they may suffer from the recommendation quality issue. In fact, taking into account just the preferences of the target user and ignoring the opinions coming from other users, the system may recommend a matching place with poor quality from the social opinion standpoint. In order to overcome this problem, a large number of works exploit the users' location histories and CF models. Leveraging on the other users' past actions (e.g., reviews or check-ins), these location-based RS are able to improve the quality of the recommendations by ignoring poorly-reviewed locations that otherwise match the target user's profile. As we argued in Section 3.1, the crucial point in the CF's recommendation process is represented by the selection of neighbors users (or items). The early solutions define their neighbors selection according to the first law of geography of Tobler [121], which claims that "everything is related to everything else, but near things are more related than distant things". For instance, the work in [122] considers only individuals who live near the location from which the target user has made the query, while in [123] candidate POIs are pruned by using a spatial range around the user's current location. In addition to the spatial distance between places and users, several works propose to use a different kind of information in order to further improve the quality of the recommendations. For example, Shi et al. [124] proposes a category-regularized matrix factorization model, taking into account the category of the places visited by the user in the past (e.g., "Italian restaurant" or "Irish pub").

Sentiment analysis can also be used in order to build a more fine-grained model of the user's preferences [125,126]. Yang et al. [127] exploit the user's comments left in the check-ins in a tensor factorization framework to build a personalized location ranking system, while the system proposed in [126] considers as neighbors not just the users with similar preferences, but also those with comparable needs. Specifically, it clusters the places' reviews according to different aspects, such as the topics covered in the text (e.g., service, food, facilities), the intent and the nationality of the author. In a recent work, Wang et al. [128] face the problem of POIs recommendation considering the user interest drift [129]. According to this phenomenon, users tend to have different interests when they travel out of their hometown. To this aim, the authors propose a model to simulate the decision-making process of users' check-in behaviors, both in hometown and out-of-town areas, using both geographical clustering and topic modeling (i.e., Latent Dirichlet Allocation – LDA [130]) on users' check-ins and reviews.

Furthermore, other works explore the influence of social aspects in location-based recommendations. Assuming that friends share many more preferences than strangers, Ye et al. [131] limit the choice of neighbors to the user's friends only. They extend their work in [132] by selecting candidate neighbors taking into account the user preferences extracted from check-ins histories and social connections, but also the geographical distance between users and candidate places. They also perform a comparative evaluation among different CF and random walk-based approaches, proving that the geographical factor has more impact on the accuracy of the recommendations than personal social relations. However, social links better address the cold-start problem because social friends may provide potentially relevant POIs to new users without location histories. Therefore, the combination of different aspects, such as geographical influence, social relations and user preferences, ensures to obtain the highest performance in recommending individual locations [133].

As previously described, SRS can provides people/friend recommendations mainly exploiting the underlying social structure and user interaction patterns [74,77,134]. However, it has been proven that location has significant correlation with users' social behaviors in their real life, such as friendship relationships [135,136]. Therefore, LBSNs can provide also a new way to make people/friend recommendations taking into account the users' location histories. For example, Zheng et al. [137] and Ying et al. [138] focus on finding popular users in LBSNs. They consider as "popular" the users with more knowledge about the locations (e.g., those who made several check-ins) and they assume that this kind of users are able to provide high quality location recommendations. A large body of research (e.g., [139-143]) leverages on users' location histories to improve the effectiveness of friend recommendations. Here, the main idea is that the history of visited places can reveal preferences, and thus people with similar location histories may have

similar preferences and may become friends more likely. Scellato et al. [143] propose a supervised learning framework that exploits users' visited locations to predict new links among *place-friends*, i.e., users who visit the same places, and *friends-of-friends*. Their results show how the inclusion of information about places, and related user activities, offers high link prediction probability for the friend recommendation task. Another example is represented by Xiao et al. [140], which estimates the similarity between users according to the semantic of their visited locations (e.g., restaurants or cinemas). This allows the system to recommend connections between users who have different geographic behaviors (e.g., living in different cities), but share similar semantic behaviors (e.g., they frequently visit the same types of locations).

Leveraging on the data shared by users in LBSNs, it is also possible to extract useful knowledge about locations and user activities. For instance, users may be interested in knowing which activities (e.g., dining, shopping, watching movies/shows, enjoying sports/exercises) can be practiced in a given place. Activity recommendations can be performed by exploiting both the meta-data associated with the locations (e.g., tags) [144] and the users' location histories [145]. On the other hand, other works provide activity recommendations by exploiting information from all the users with collaborative-based approaches. For instance, Zheng et al. [146] and Symeonidis et al. [147] propose the use of a 3-order tensor to characterize each user check-in operation. In this case, the tensor includes the current user's location, the performed activity and a rating on that activity. They use factorization methods on these tensors in order to provide location-specific activity recommendations.

The last type of recommendations that can benefit from the use of geographical information is the content recommendation, with specific reference to social media contents. In this case, the system provides users with suggestions of photos, videos, or other web content they might like, which have been shared by other users through a OSN. Several works exploit the histories of users location to improve the quality of social media recommendations. For instance, Kawakubo and Yanai [148] proposes a picture ranking algorithm that exploits the locations in order to improve the relevance of search results. Silva and Martins [149] presents a method to suggest the most relevant tags for georeferenced photos, and [150] builds a location-based RS aimed at increasing the diversity of recommendations, while mitigating popularity bias of web contents. In a recent work, Pálovics et al. [151] address the problem of recommending highly volatile items, as locations or events characterized by strict time constraints (e.g., a concert or an expo). Specifically, they use both online machine learning and matrix factorization techniques to recommend to their users the most relevant hashtags for a given geo-referenced tweet at a given time. Their solution performed much better than online matrix factorization [152] and content-based [153] methods, which justifies the importance of temporal and geographic information in the socialmedia recommendation task.

5. Recommender Systems for Mobile Social Networks

In Mobile Social Networks (MSN), typically, a mobile device should be able to autonomously share information with other devices in proximity in order to discover useful contents for its user. In this context, RS represent a useful tool to improve the content dissemination in MSN, by proactively suggesting to the users the information discovered in the nearby. However, given the distributed nature of MSN, the solutions proposed for OSN (discussed in Section 4) cannot be simply extended to this new scenario. In fact, RS for OSN generally rely on standard client/server centralized models, where the recommendation engine runs on the server side (or a cloud-based infrastructure) and processes the requests coming from fixed and mobile clients. In addition, the recommendation task in MSN significantly differs from that defined for the online environment. As discussed in Section 2, RS for OSN typically focus on learning a predicting model for missing values in the ratings matrix. This matrix can be seen as a systems' global knowledge about both the available items and users' preferences. Thus, the recommendation engine learns a single global model which will be used to perform the recommendations for all its users. On the contrary, in MSN each device may be aware of just a small part of the global information. This local knowledge is initially related to information about the local user only, then it grows up with those exchanged with encountered users and devices, through D2D communications. Therefore, in MSN each device should learn a recommendation model focused on the preferences of its local user.

The first approaches proposed in the literature for RS in distributed environments refer to the Peer-to-Peer (P2P) paradigm. In a P2P architecture, RS are fully distributed over several nodes, which act as client and server at the same time [154]. The solutions proposed in this context (e.g., [155–157]) are typically based on Neighborhood CF. Specifically, the user–item ratings are maintained in a distributed way on the P2P infrastructure and, when a peer has to make a prediction for its local user, it exploits P2P lookup methods to find and retrieve the relevant ratings necessary to locally calculate the rating for the target item.

By shifting the recommendation process to the client side, P2P solutions are able to alleviate some critical points of centralized methods; for example, the possible bottleneck represented by the server, or security and privacy issues that can arise when data is handled by a single organization [158,159]. However, by using standard P2P systems, peers should be able to communicate to each other by using infrastructured networks (i.e., accessing the Internet), in order to have a constant access to the entire knowledge of the network. This is not always possible in MSN and, in the last few years, some RS explicitly tailored for MSN and opportunistic networks have been presented in the literature.

We classify the proposed solutions based on the RS method used: (i) CF, and (ii) tag-based, previously described in Sections 3.1 and 4.2, respectively. Hereafter we present the main solutions belonging to the two categories.

5.1. Collaborative Filtering

The first set of RS proposed in the field of MSN uses CF. In this case, the recommendation task is performed on each local device, and it is based on the users' ratings exchanged among nodes via D2D communications.

The proposed solutions differ in two main aspects: (i) the metrics used to measure the similarity between users, and/or (ii) the heuristics implemented to reduce the computational complexity of CF in a mobile environment. The simplest similarity measures refer to the rating-based metrics used in standard CF (e.g., the Pearson Coefficient discussed in Eq. (9)), while other solutions take into account additional context information to further improve the quality of the recommendations.

5.1.1. OppCF

De Spindler et al. [160] are among the first in the literature who investigated the use of RS in opportunistic environments. Specifically, they use a distributed user-based CF approach to predict ratings based on shared information among co-located mobile devices. They present a preliminary idea of "social context" shared among mobile devices during their opportunistic contacts and aimed at improving the recommendation process. In this case, the social context is identified by the co-location of users related to a specific event (e.g., attending a conference or a concert). This 90

condition is the basic assumption of the authors to identify similar users, as those who are physically co-located in a specific time interval. The authors do not consider the possibility to exploit also the occasional co-location of users as potential source of interesting contents. This pioneering solution has never been evaluated, thus its accuracy cannot be assessed.

5.1.2. MobHinter

MobHinter [161] presents two main contributions: (i) the definition of Affinity networks aimed at identifying similar users, and (ii) an epidemic protocol designed to distribute and exchange users' ratings among mobile devices in ad-hoc networks. The authors focus on the concept of affinity on ratings, by calculating the percentage of similar ratings provided by two users on a common set of items, and by assigning a threshold value to identify similar users. Therefore, to apply CF method based on users' affinities, MobHinter proposes an epidemic protocol to exchange a set of ratings information through ad hoc communications: (i) the list of ratings of the local node, (ii) the list of ratings of its current neighbors and (iii) the ratings of the nodes encountered in the past. Then, each node is able to locally compute its affinity network, and to apply standard CF methods to predict missing ratings and provide the recommendation. The authors proposed also optimized versions of the protocol, by exchanging only the list of ratings of the local node or by excluding those of the nodes encountered in the past.

Authors evaluate their proposal using the MovieLens dataset, by simulating a fully connected scenario (to estimate the optimal behavior) and an ad hoc network scenario based on a random mobility model. In the former scenario, they compared their model with other user-based CF methods, evaluating the prediction accuracy (i.e., MAE and RMSE). Then, they performed a simulation to evaluate the number of random meetings needed on average by users to approximate the reference prediction accuracy obtained in the fully connected scenario. The authors showed that by exchanging all the ratings information known by each user, the system approximates the reference accuracy with a limited number of meetings. However, the use of a random mobility model is not realistic to simulate an opportunistic networking scenario as largely demonstrated in the literature [162].

5.1.3. DiffeRS

Del Prete and Capra [163] propose an interesting approach to reduce the complexity of CF in MSNs. They propose diffeRS [163], a decentralized RS aimed at locally classify the user as mass-like minded or individual, based on the increasing knowledge about the other encountered users and their profiles. In fact, also in this case, mobile devices exchange the users profile through D2D communications (by exploiting Bluetooth technology, in this case). For each recommendation task, diffeRS evaluates the average deviation of the local user's profile with respect to the preferences expressed by the local community as a whole (i.e., the average ratings). If the local user is identified as mass-like minded, the recommendations are simply based on the average of the community preferences, thus reducing at the minimum the computation on the mobile device. On the other hand, if the local user is an individual, a user-based CF approach is applied, considering only other individual users that differ from the community's preferences in the same proportion of the local user. To this aim, the system selects only those neighbors that share with the local user at least one rated item, and it computes the difference among their co-ratings to define a measure of their similarity with respect to their deviation from the average of the community preferences. In this way, diffeRS further reduces the sparsity of the rating matrix used in CF and, consequently, the recommendation complexity for both user profiles, by exchanging a limited amount of information during the opportunistic contacts. The authors evaluate diffeRS both in a fully connected scenario and in an opportunistic one by using two datasets: MovieLens to model users' ratings, and MIT Reality Mining [164] to model users' mobility. In the fully connected scenario, the authors evaluate the prediction accuracy in terms of MAE and they demonstrate that the proposed correlation measure, divided between *mass-like minded* and *individual*, performs better than the standard Pearson correlation.

In the opportunistic scenario, they randomly map100 users from the MovieLens ratings dataset to the 100 users of the MIT Reality Mining dataset. Predictions are computed at regular intervals of time, comparing accuracy (i.e., MAE) and coverage of diffeRS with respect to CF based on Pearson correlation. The obtained results show that diffeRS always outperforms the user-based CF.

5.1.4. LocPref

In a recent work, Zhao et al. [165] propose the use of CF to recommend privacy preferences depending on the user location. They refer to a Location Sharing Service (LSS) scenario, where users can share their current location using the "check-in" functionality of social network applications. The declaration of a user to be in some specific location (e.g., a clinic or religious sites) may be considered as a sensitive information, and people usually may not wish to share them with the others [166]. LSS can autonomously share location information, following the general privacy settings of the user. LocPref is aimed at recommending the user with specific privacy settings for sharing her current location through the LSS. To this aim, it leverages on the assumption that people who visited the same location can share similar privacy preferences. Therefore, users' mobile devices can exchange location-privacy preferences through opportunistic communications and locally compute their recommendations.

The authors also propose a simple but elegant mechanism to reduce the vulnerability of CF to *shilling attack* [167]: the tentative of malicious users (i.e., attackers) to bias the recommendations by giving false information and/or ratings [168]. The attackers can create fake profiles (also called *shill profiles*) in order to inject modified data into the recommendation process. For instance, in the context of LSSs, a business owner may force the customers to share their location when they visit her shops in order to increment their visibility on social media. In the same way, she may want to prevent the customers' check-ins at competitors' shops to reduce their popularity.

In order to prevent shilling attacks, authors propose a reputation scheme based on the encounter frequency of nodes. Every time a node receives a user's profile from another node, it increments a counter associated with that specific profile and considers it as the reputation level of the encountered user. Once the system needs to perform a recommendation, it uses only those profiles whose counter is greater than the average reputation of all the profiles in the local cache (i.e., the average number of encounters of all nodes in the cache). An attacker usually creates multiple fake profiles, but in an opportunistic environment, it can share just one of them during a contact. For this reason, this simple mechanism is able to associate a lower reputation with the shill profiles than those calculated for real ones.

To evaluate the proposed solution, locPref implements two different schemes for the opportunistic data exchange: (i) nodes exchange only the local user profile (in terms of privacy settings) at each opportunistic contact (*D-Ind*); and (ii) nodes exchange both the local profile and those previously received by the others (*D-Set*). In order to obtain a reference evaluation metric, the authors simulate the system in a fully connected environment and then in a mobile scenario. They exploit the *st_andrews/locshare* dataset available in CRAWDAD [169], which contains the location-privacy preferences of 40 users collected in the city of St Andrews (UK), and they simulate the user mobility patterns based on St Andrews road map and selected POIs in the city.

Simulations results show a comparison between *D-Ind* and *D-Set* versions of locPref, highlighting the time needed to reach the same accuracy of the fully connected scenario (comparable with a centralized solution) and the advantages of exploiting the exchange of additional profile information to compute the recommendations. In addition, the authors demonstrate the efficiency of their reputation-based mechanism, which allows to prevent the Recommender System from being abused by shilling attacks.

5.2. Exploiting tags

As we described in Section 4.2, exploiting user-defined tags (i.e., folksonomies) is an effective method to improve the quality of the recommendations. Tags can be also used in MSN to further characterize contents even from a semantic point of view. However, in order to implement tag-based MSN solutions, we need a mechanism which is able to infer the relations among different tags using just the user's preferences and the limited information obtained by other nodes in proximity. In the following, we present in detail the two solutions proposed in the literature, which exploit tags in MSNs: ICe-Habit [170] and p-PLIERS [112].

5.2.1. Information-Centric Habit (ICe-Habit)

To the best of authors' knowledge, ICe-Habit [170] is the first approach proposed in the literature that exploits folksonomies to improve the performance of a content dissemination protocol (i.e., Habit [171]) for opportunistic networks. In ICe-Habit, each user's profile is characterized by a vector of tags used in the past by the user to describe her contents. User profiles are exchanged by nodes during opportunistic contacts. In this way, each node can locally build a tag co-occurrence matrix *M*. Each entry of the matrix, M[i, j] represents the number of items characterized by the co-occurrence of tags (t_i, t_j) . Therefore, *M* is used to select contents to be forwarded to potentially interested users in opportunistic networks.

Whenever a user creates a new content, she associates a set of tags $T = \{t_1, \ldots, t_n\}$ with it. Then, the system performs the *tagexpansion* procedure to further enrich the set of tags *T*. Specifically, the system queries the matrix *M* to extract the *k* tags (i.e., T_k) with the highest value of co-occurrence with those specified in *T*. Then, the expansion set of tags $T' = T \cup T_k$ is associated to the new content. Finally, the item is opportunistically forwarded to the potentially interested users, i.e., each user *x* whose profile T_x contains at least one of the tags in *T'*.

The authors evaluate ICe-Habit, via simulation, in an opportunistic scenario. To this aim, they use the MIT Reality Mining datasets, which contains the contact traces of 96 nodes. In addition, the MovieLens dataset has been used to create a user's profile for each node. Specifically, each profile contains the set of tags the user has created to tag movies in the dataset. Furthermore, for each $\langle user_i, tag_i \rangle$ pair, the MovieLens dataset contains also the timestamp information in which the tag_i has been created by the *user_i*. The authors use the timestamp information to map the creation of new items on the simulated time. In order to measure the performance of the proposed solution, authors use the following approach: when a node creates a new content, it does not include all the tags it would normally associate with it, but it randomly drops 50% of the item's tags. In this way, authors measure the Tag Recovery and Destination Recovery, which are, respectively, the percentage of the dropped tags and correct destination reached by each message that could be recovered by the tag expansion method.

Although the results demonstrate that the tag expansion improves the content dissemination's performance in a MSN, they also prove that this approach is not able to correctly infer the semantic relations between different tags. In fact, considering just the tags' co-occurrences, nodes could receive more items than those really interesting for them because users may not be interested in the topics represented by the expanded tags.

5.2.2. Pervasive PLIERS (p-PLIERS)

Recently, we proposed a novel framework for the content discovery and evaluation in MSN called *pervasive PopuLarity-based ItEm Recommender System* (*p-PLIERS*) [112]. It can be used both as standalone RS for MSN, and as a useful support tool for content dissemination or routing protocols in fully distributed environments (e.g., opportunistic networks).

Differently from ICe-Habit [170], p-PLIERS does not leverage just on the tags' co-occurrences. It exploits all the folksonomy's information, i.e., the relations among users, shared or created items, and tags used to semantically describe the contents. This information is locally maintained by each node, and modeled as a tripartite users-items-tags graph, which is called *Local Knowledge Graph* (*LKG*). The LKG of each node merges the information about the items created or downloaded by its local user, with the local knowledge of other encountered nodes, obtained through opportunistic communications. Specifically, when a node creates a new item, it updates its LKG by inserting the relation between its user entity, the generated item and the related tags. Similarly, when a node encounters another node in the network, the two nodes exchange their LKGs, and locally integrate them with the received information.

Once each node has updated its local LKG with the information coming from the other node, p-PLIERS implements PLIERS [111] (described in Section 4.2) to evaluate the relevance of the new available items with respect to the interests of its local user.

As a first set of experiments, we evaluate the accuracy of PLIERS recommendations in a centralized scenario with respect to other solutions proposed for MSN: the Tag Expansion mechanism exploited by ICe-Habit [170] and user-based CF used in similar works (e.g., MobHinter [161] or diffeRs [163]). To this aim, we extract a real dataset from Twitter using the Twitter Streaming API⁷. Specifically, we have downloaded the tweets generated during a big event (i.e., the World Food Day at Expo 2015) in the urban area of Milan, and we have built a tripartite graph composed by more than 5000 users, 2946 tweets, and 3202 hashtags. This represents a realistic folksonomy user-item-tag graph of online tagged contents related to a popular event. To assess the accuracy of the recommendations, we perform a link prediction task on a tripartite graph by removing one link from each user connected at least to 5 items with popularity greater than one (i.e., the number of connected users). Then, we calculate the performance of PLIERS, Tag Expansion and user-based CF in terms of precision and recall, based on the number of links that are included in the recommendations of each algorithm (i.e., the "recovered links"). In the centralized scenario, PLIERS outperforms the two reference algorithms for both considered measures. Specifically, it reaches a precision score up to eight times higher than that obtained by Tag Expansion and a score 40% higher than CF. With regard to the Recall measure, PLIERS obtains a score around 50% higher than Tag Expansion and around 30% higher than CF. These results indicate that PLIERS, exploiting all the information contained in the folksonomy graph (i.e., user-item-tag relationships), is able to obtain better results than the other solutions proposed for MSN, which consider only part of the user-item or item-tag relationships.

In order to validate the overall framework, we have simulated three realistic scenarios: a big event (Expo 2015), a conference

⁷ https://developer.twitter.com/en/docs/tweets/filter-realtime



Fig. 13. Map of Expo 2015 area with the position of five of the simulated communities [112].

venue (ACM KDD 2015), and a working day in the Helsinki city center. In this case, we investigate the efficiency of p-PLIERS in a fully decentralized environment, where the contents are dynamically generated over time by nodes. In particular, we measure the similarity of the recommendations generated by the framework on the nodes' LKGs and on the global knowledge graph (GKG), that is, the union of all the LKGs at each time-step of the simulations. To this aim, we use both real and synthetic mobility traces, and we build three different datasets from Twitter for the nodes' content generation during the simulations. Specifically, for the Expo 2015 scenario, we simulate different sets of nodes (i.e., 250, 500, and 900 nodes) moving in the Expo area of Milan using the HCMM [172] human mobility model. Using this model, we are able to simulate the presence of different communities, which fits well with the reference scenario, as the Expo area was divided into several pavilions. Fig. 13 depicts the simulated area in the Expo scenario with the communities of nodes used in HCMM. In addition, we use the dataset related to the World Food Day in Expo 2015 as content (i.e., tweets) generation pattern of the nodes during the simulations.

As a second dynamic scenario, we consider a school campus during a conference event, where people stay most of the time within rooms, but they regularly gather at breaks (e.g., coffee or lunch breaks). Therefore, we use real contact traces representing the physical interactions of a group of students, professors, and staff of an American high school during a typical school day (i.e., the high resolution Human Contact Network - HCN [173]). As American high schools are not organized into classes as in Europe, but rather around study tracks, and students are free to decide which lectures to attend, we think that their movements can be assimilated to those of people attending a large conference. For this reason, we have downloaded the tweets generated during a large computer science conference (i.e., the 21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining), and we use this data to simulate the nodes' content generation during the simulation.

Finally, in order to simulate the use of p-PLIERS on a larger scale than the previous scenarios, we extract the contact traces of a typical working day in Helsinki using a realistic human mobility model highly customized on the considered area. Specifically, we have used the ONE [174] simulator, which implements the *Working Day Mobility Model (WDMM)* [175] and the map of the city center of Helsinki. Then, we have downloaded the tweets generated within the same geographic area to use them as content generated by nodes during the simulation.

In the best case, i.e., the conference scenario, the average similarity reaches 80% already in two hours simulated time, over eight hours duration. In the worst case, i.e., Helsinki scenario, the average similarity reaches 80% at the end of the simulation, due to the reduced number of contacts during the working day.

This implies that, in dense scenarios, nodes require just few contacts to well approximate the global knowledge with their local graphs, demonostrating that p-PLIERS is able to provide effective recommendations, comparable to those achievable if global knowledge is accessible. For the urban scenario of Helsinki, in which the considered area is much larger than that of the other scenarios and the density of nodes is lower, the results indicate that a larger time window is required for a good approximation of the global knowledge about contents in the network. With a view to smart cities, a possible solution to improve the diffusion of knowledge and the accuracy of p-PLIERS might be based on exploiting also the public transportation system nodes (e.g., buses, trams, or taxis) as additional information carriers.

In Table 1 we summarize and compare the solutions presented in the previous sections specifying, for each of them, the method implemented for recommendation, the context information used to optimize the process, and the details of their performance evaluation in terms of (i) dataset, (ii) mobility traces and (iii) performance metrics.

6. Concluding remarks and open challenges

In this work, we have presented a survey of the main RS proposed in the literature for Online and Mobile Social Networks, with particular attention to the use of social context information to improve the recommendation process. We describe advantages and drawbacks of standard recommendation techniques in these environments and we highlight RS challenges in the fully distributed environment of Mobile Social Networks. Several solutions have been proposed for OSN, while the study of efficient RS for MSN is still in its infancy and presents several open research challenges, like the use of additional context information, specifically related to the mobile environment, to further optimize the recommendation process.

Due to their sensing capabilities and the constant presence in human daily life, modern personal devices (e.g., smartphones and tablets) represent the bridge between the cyber and physical worlds, characterizing the situation in which the user is involved during everyday life. Therefore, mining the information provided by these devices, and properly combining them with OSN data, can provide a more accurate model of the user's context and preferences. For example, the social context of the user may be defined combining the virtual social relationships extracted from OSNs with the physical contacts among the devices and the information contained in the user's personal device (e.g., the contacts saved in the address book, call logs, and messages). In this case, the user's preferences and needs can be inferred from Table 1 RS for MSN.

Algorithm	RS method	Context info	Dataset	Evaluation mobility scenario	Metrics
OppCF [160]	User-based CF	Location Time Ratings			
Mobhinter [161]	User-based CF	Ratings	MovieLens	Random	RMSE MAE
DiffeRS [163]	User-based CF + heuristic	Ratings	MovieLens	MIT Reality Mining [164]	MAE Coverage
LocPref [165]	User-based CF + heuristic	Location Time Ratings	LocShare	Map-based	Custom Accuracy metric
ICe-Habit [170]	Tags correlations	Tags	MovieLens	MIT Reality Mining [164]	Recall
p-PLIERS [112]	Graph- and diffusion-based	Users Items Tags	Twitter	HCMM [172] HCN [173] WDMM [175]	Precision Recall

heterogeneous sources of data, e.g., the web browser history, the actions performed in OSNs, the generated contents both in OSN and MSN, the visited locations, and the set of most used mobile apps.

The availability of this heterogeneous information represents a crucial aspect in building and evaluating new recommending solutions. However, in order to validate and evaluate the possible solutions in realistic environments, it is necessary to generate new datasets, collecting all the contextual information that can characherizes characterise the mobile environment and can be exploited in MSN. Collecting new datasets is a very time consuming task, requiring the implementation of a real prototype application and the involvement of a high number of mobile users in real environments. A possible solution is represented by the use of synthetic (but realistic) context data generator (e.g., [176,177]) but, at the moment, none of these solutions is able to integrate realistic temporal settings in the data generation process. Since user's preferences are highly dynamic, especially in a mobile environment, the use of an obsolete preference model could negatively affect the quality of the recommendations.

In addition, several frameworks have been recently proposed to evaluate RS performance in OSN, e.g., *LibRec* [178], *LensKit* [179], *RankSys* [180], and *CARSKit* [181]. They provide an implementation of some RS methods (e.g., CF, MF) and performance metrics, but they are developed exclusively for centralized RS, not considering the unique characteristics of MSNs (e.g., the limited knowledge available in each device). Therefore, in order to perform a fair comparison among different RS for MSN, we should define a common evaluation framework with the following three main characteristics:

- Mobility: RS should be evaluated in realistic mobile scenarios. Therefore, the evaluation framework should include datasets of real mobility traces and human mobility models.
- Data: as we previously pointed out, RS for MSN should characterize the user's context with several heterogeneous data. To this aim, the framework should also include proper datasets or context data generators. In addition, given the extremely dynamic nature of the MSN environment, during the simulation nodes should be able to generate data over time.
- *Evaluation*: the framework should include the main evaluation metrics proposed in the literature for the two main recommendation tasks (i.e., prediction and ranking). Furthermore, due to the typically limited duration of the opportunistic contacts, it is also needed to consider the time complexity of the RS in MSN.

In fact, if the RS is not able to provide recommendations in a limited time window, a node may miss the opportunity to exchange useful data with another one in proximity.

All these aspects represent a wide research area, including not only the definition of new RS algorithms, but also the creation of experimental testbeds including new opportunistic networking protocols, opportunistic sensing features and context reasoning tools.

Acknowledgments

This work was carried out in the framework of the INTESA project (CUP CIPE D78I16000010008), co-funded by the Tuscany Region (Italy) under the Regional Implementation Programme for Underutilized Areas Fund (PAR FAS 2007–2013) and the Research Facilitation Fund (FAR) of the Ministry of Education, University and Research (MIUR).

References

- F. Ricci, Mobile recommender systems, Inf. Technol. Tour. 12 (3) (2010) 205–231.
- [2] I. Guy, D. Carmel, Social recommender systems, in: Proceedings of the Twentieth International Conference Companion on World Wide Web, WWW '11, ACM, New York, NY, USA, 2011, pp. 283–284, doi:10.1145/1963192.1963312.
- [3] M. Conti, S. Giordano, M. May, A. Passarella, From opportunistic networks to opportunistic computing, IEEE Commun. Mag. 48 (9) (2010) 126–139, doi:10. 1109/MCOM.2010.5560597.
- [4] M. Conti, A. Passarella, S.K. Das, The Internet of People (IoP): A new wave in pervasive mobile computing, Pervasive Mob. Comput. 41 (2017) 1–27.
- [5] Y. Shi, M. Larson, A. Hanjalić, Collaborative filtering beyond the user-item matrix: a survey of the state of the art and future challenges, ACM Comput. Surv. (CSUR) 47 (1) (2014) 3.
- [6] Y.-J. Park, A. Tuzhilin, The long tail of recommender systems and how to leverage it, in: Proceedings of the 2008 ACM Conference on Recommender Systems, ACM, 2008, pp. 11–18.
- [7] T. Mahmood, F. Ricci, Learning and adaptivity in interactive recommender systems, in: Proceedings of the Ninth International Conference on Electronic Commerce, ACM, 2007, pp. 75–84.
- [8] T. Mahmood, F. Ricci, Improving recommender systems with adaptive conversational strategies, in: Proceedings of the Twentieth ACM Conference on Hypertext and Hypermedia, ACM, 2009, pp. 73–82.
- [9] N. Taghipour, A. Kardan, S.S. Ghidary, Usage-based web recommendations: a reinforcement learning approach, in: Proceedings of the 2007 ACM Conference on Recommender Systems, ACM, 2007, pp. 113–120.
- [10] C.J. Willmott, K. Matsuura, Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance, Clim. Res. 30 (1) (2005) 79–82.
- [11] L. Baltrunas, T. Makcinskas, F. Ricci, Group recommendations with rank aggregation and collaborative filtering, in: Proceedings of the Fourth ACM Conference on Recommender Systems, RecSys '10, ACM, New York, NY, USA, 2010, pp. 119–126, doi:10.1145/1864708.1864733.

- [12] S. Balakrishnan, S. Chopra, Collaborative ranking, in: Proceedings of the Fifth ACM International Conference on Web Search and Data Mining, ACM, 2012, pp. 143–152.
- [13] J.L. Herlocker, J.A. Konstan, L.G. Terveen, J.T. Riedl, Evaluating collaborative filtering recommender systems, ACM Trans. Inf. Syst. (TOIS) 22 (1) (2004) 5–53.
- [14] A. Gunawardana, G. Shani, Evaluating recommender systems, in: Recommender Systems Handbook, Springer, 2015, pp. 265–308.
- [15] F. Ricci, L. Rokach, B. Shapira, Introduction to recommender systems handbook, in: Recommender Systems Handbook, Springer, 2011, pp. 1–35.
- [16] J. Tang, X. Hu, H. Liu, Social recommendation: a review, Soc. Netw. Anal. Min. 3 (4) (2013) 1113–1133.
- [17] X. Ning, C. Desrosiers, G. Karypis, A comprehensive survey of neighborhood-based recommendation methods, in: Recommender Systems Handbook, Springer, 2015, pp. 37–76.
- [18] J. Herlocker, J.A. Konstan, J. Riedl, An empirical analysis of design choices in neighborhood-based collaborative filtering algorithms, Inf. Retr. 5 (4) (2002) 287–310.
- [19] C.C. Aggarwal, Recommender Systems, Springer, 2016.
- [20] B. Mobasher, H. Dai, T. Luo, M. Nakagawa, Effective personalization based on association rule discovery from web usage data, in: Proceedings of the Third International Workshop on Web Information and Data Management, ACM, 2001, pp. 9–15.
- [21] M.-L. Shyu, C. Haruechaiyasak, S.-C. Chen, N. Zhao, Collaborative filtering by mining association rules from user access sequences, in: Proceedings of the 2005 International Workshop on Challenges in Web Information Retrieval and Integration, WIRI'05, IEEE, 2005, pp. 128–135.
- [22] K. Miyahara, M. Pazzani, Collaborative filtering with the simple Bayesian classifier, Proceedings of the Eighteenth Pacific Rim International Conference on PRICAI 2000 Topics in Artificial Intelligence, 2000, pp. 679–689.
- [23] X. Su, T.M. Khoshgoftaar, Collaborative filtering for multi-class data using belief nets algorithms, in: Proceedings of the Eighteenth IEEE International Conference on International Conference on Tools with Artificial Intelligence, IC-TAI'06, IEEE, 2006, pp. 497–504.
- [24] Z. Xia, Y. Dong, G. Xing, Support vector machines for collaborative filtering, in: Proceedings of the Forty-fourth Annual Southeast Regional Conference, ACM, 2006, pp. 169–174.
- [25] R. Salakhutdinov, A. Mnih, G. Hinton, Restricted Boltzmann machines for collaborative filtering, in: Proceedings of the Twenty-fourth International Conference on Machine Learning, ACM, 2007, pp. 791–798.
- [26] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, T.-S. Chua, Neural collaborative filtering, in: Proceedings of the Twenty-sixth International Conference on World Wide Web, International World Wide Web Conferences Steering Committee, 2017, pp. 173–182.
- [27] Y. Koren, R. Bell, C. Volinsky, Matrix factorization techniques forrecommender systems, Computer 42 (8) (2009) 30–37.
- [28] M. Pazzani, D. Billsus, Content-Based Recommendation Systems, in: P. Brusilovsky, A. Kobsa, W. Nejdl (Eds.), The Adaptive Web: Methods and Strategies of Web Personalization, Springer Berlin Heidelberg, Berlin, Heidelberg, 2007, pp. 325–341.
- [29] P. Lops, M. De Gemmis, G. Semeraro, Content-based recommender systems: state of the art and trends, in: Recommender Systems Handbook, Springer, 2011, pp. 73–105.
- [30] T. Miranda, M. Claypool, A. Gokhale, T. Mir, P. Murnikov, D. Netes, M. Sartin, Combining content-based and collaborative filters in an online newspaper, in: Proceedings of 1999 ACM SIGIR Workshop on Recommender Systems, Citeseer, 1999.
- [31] G. Adomavicius, A. Tuzhilin, Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions, IEEE Trans. Knowl. Data Eng. 17 (6) (2005) 734–749.
- [32] F. Yu, A. Zeng, S. Gillard, M. Medo, Network-based recommendation algorithms: a review, Phys. A Stat. Mech. Appl. 452 (2016) 192–208.
- [33] L. Page, S. Brin, R. Motwani, T. Winograd, The PageRank Citation Ranking: Bringing Order to the Web., Technical report, Stanford InfoLab, 1999.
- [34] T.H. Haveliwala, Topic-sensitive PageRank, in: Proceedings of the Eleventh International Conference on World Wide Web, ACM, 2002, pp. 517–526.
- [35] A. Hotho, R. Jäschke, C. Schmitz, G. Stumme, FolkRank: a ranking algorithm for folksonomies, in: Proceedings of the 2006 Conference on Live Working Association (LWA), vol. 1, 2006, pp. 111–114.
- [36] M.-H. Tsai, C. Aggarwal, T. Huang, Ranking in heterogeneous social media, in: Proceedings of the Seventh ACM International Conference on Web Search and Data Mining, ACM, 2014, pp. 613–622.
- [37] M. Gori, A. Pucci, V. Roma, I. Siena, Itemrank: a random-walk based scoring algorithm for recommender engines., in: Proceedings of the 2007 International Joint Conference on Artificial Intelligence, IJCAI, vol. 7, 2007, pp. 2766–2771.
- [38] K. Lakiotaki, N.F. Matsatsinis, A. Tsoukias, Multicriteria user modeling in recommender systems, IEEE Intell. Syst. 26 (2) (2011) 64–76.
- [39] F. Fouss, A. Pirotte, J.-M. Renders, M. Saerens, Random-walk computation of similarities between nodes of a graph with application to collaborative recommendation, IEEE Trans. Knowl. Data Eng. 19 (3) (2007) 355–369.
- [40] H. Yildirim, M.S. Krishnamoorthy, A random walk method for alleviating the sparsity problem in collaborative filtering, in: Proceedings of the 2008 ACM Conference on Recommender Systems, ACM, 2008, pp. 131–138.
- [41] T. Zhou, Z. Kuscsik, J.-G. Liu, M. Medo, J.R. Wakeling, Y.-C. Zhang, Solving the apparent diversity-accuracy dilemma of recommender systems, Proc. Natl. Acad. Sci. 107 (10) (2010) 4511–4515.

- [42] T. Zhou, J. Ren, M. Medo, Y.-C. Zhang, Bipartite network projection and personal recommendation, Phys. Rev. E 76 (4) (2007) 046115.
- [43] J.-G. Liu, T. Zhou, Q. Guo, Information filtering via biased heat conduction, Phys. Rev. E 84 (3) (2011a) 037101.
- [44] J.-G. Liu, Q. Guo, Y.-C. Zhang, Information filtering via weighted heat conduction algorithm, Phys. A Stat. Mech. Appl. 390 (12) (2011b) 2414-2420.
- [45] L. Lü, W. Liu, Information filtering via preferential diffusion, Phys. Rev. E 83 (6) (2011) 066119.
- [46] T. Zhou, L.-L. Jiang, R.-Q. Su, Y.-C. Zhang, Effect of initial configuration on network-based recommendation, Europhys. Lett. (EPL) 81 (5) (2008) 58004.
- [47] C. Liu, W.-X. Zhou, Heterogeneity in initial resource configurations improves a network-based hybrid recommendation algorithm, Phys. A Stat. Mech. Appl. 391 (22) (2012) 5704–5711.
- [48] G. Adomavicius, A. Tuzhilin, Context-aware recommender systems, in: Recommender Systems Handbook, Springer, 2011, pp. 217–253.
- [49] G. Adomavicius, R. Sankaranarayanan, S. Sen, A. Tuzhilin, Incorporating contextual information in recommender systems using a multidimensional approach, ACM Trans. Inf. Syst. (TOIS) 23 (1) (2005) 103–145.
- [50] G. Adomavicius, A. Tuzhilin, Context-aware recommender systems, in: Recommender Systems Handbook, Springer, 2015, pp. 191–226.
- [51] R. Burke, Hybrid recommender systems: survey and experiments, User Model. User-Adapt. Interact. 12 (4) (2002) 331–370.
- [52] V. Codina, F. Ricci, L. Ceccaroni, Exploiting the semantic similarity of contextual situations for pre-filtering recommendation, in: Proceedings of the 2013 International Conference on User Modeling, Adaptation, and Personalization, Springer, 2013, pp. 165–177.
- [53] U. Panniello, A. Tuzhilin, M. Gorgoglione, C. Palmisano, A. Pedone, Experimental comparison of pre-vs. post-filtering approaches in context-aware recommender systems, in: Proceedings of the Third ACM Conference on Recommender Systems, ACM, 2009, pp. 265–268.
- [54] A. Karatzoglou, X. Amatriain, L. Baltrunas, N. Oliver, Multiverse recommendation: n-dimensional tensor factorization for context-aware collaborative filtering, in: Proceedings of the Fourth ACM Conference on Recommender Systems, ACM, 2010, pp. 79–86.
- [55] S. Rendle, Factorization machines, in: Proceedings of the Tenth IEEE International Conference on Data Mining (ICDM), IEEE, 2010, pp. 995–1000.
- [56] S. Rendle, Z. Gantner, C. Freudenthaler, L. Schmidt-Thieme, Fast context-aware recommendations with factorization machines, in: Proceedings of the Thirty– fourth International ACM SIGIR Conference on Research and Development in Information Retrieval, ACM, 2011, pp. 635–644.
- [57] S. Rendle, Factorization machines with libFM, ACM Trans. Intell. Syst. Technol. (TIST) 3 (3) (2012) 57.
- [58] I. Guy, Social recommender systems, in: Recommender Systems Handbook, Springer, 2015, pp. 511–543.
- [59] R. Dunbar, V. Arnaboldi, M. Conti, A. Passarella, The structure of online social networks mirrors those in the offline world, Soc. Netw. 43 (Supplement C) (2015) 39–47, doi:10.1016/j.socnet.2015.04.005. URL: http://www. sciencedirect.com/science/article/pii/S0378873315000313
- [60] V. Arnaboldi, A. Passarella, M. Conti, R.I. Dunbar, Online Social Networks Human Cognitive Constraints in Facebook and Twitter Personal Graphs, Computer Science Reviews and Trends, Elsevier, Boston, 2015.
- [61] M. McPherson, L. Smith-Lovin, J.M. Cook, Birds of a feather: homophily in social networks, Ann. Rev. Sociol. 27 (1) (2001) 415–444.
- [62] J. O'Donovan, B. Smyth, Trust in recommender systems, in: Proceedings of the Tenth International Conference on Intelligent User Interfaces, ACM, 2005, pp. 167–174.
- [63] L. Terveen, D.W. McDonald, Social matching: a framework and research agenda, ACM Trans. Comput.-Hum. Interact. (TOCHI) 12 (3) (2005) 401– 434.
- [64] D. Liben-Nowell, J. Kleinberg, The link-prediction problem for social networks, J. Assoc. Inf. Sci. Technol. 58 (7) (2007) 1019–1031.
- [65] V. Martínez, F. Berzal, J.-C. Cubero, A survey of link prediction in complex networks, ACM Comput. Surv. (CSUR) 49 (4) (2016) 69.
- [66] G. Jeh, J. Widom, SimRank: a measure of structural-context similarity, in: Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2002, pp. 538–543.
- [67] I. Guy, M. Jacovi, E. Shahar, N. Meshulam, V. Soroka, S. Farrell, Harvesting with sonar: the value of aggregating social network information, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM, 2008, pp. 1017–1026.
- [68] I. Guy, M. Jacovi, N. Meshulam, I. Ronen, E. Shahar, Public vs. private: comparing public social network information with email, in: Proceedings of the 2008 ACM Conference on Computer Supported Cooperative Work, ACM, 2008, pp. 393–402.
- [69] I. Guy, I. Ronen, E. Wilcox, Do you know?: recommending people to invite into your social network, in: Proceedings of the Fourteenth International Conference on Intelligent User Interfaces, ACM, 2009, pp. 77–86.
- [70] R. Dunbar, The social brain hypothesis, Brain 9 (10) (1998) 178–190.
- [71] R.A. Hill, R.I. Dunbar, Social network size in humans, Hum. Nat. 14 (1) (2003) 53-72.
- [72] V. Arnaboldi, M. Conti, A. Passarella, F. Pezzoni, Analysis of ego network structure in online social networks, in: Proceedings of the 2012 International Conference on Privacy, Security, Risk and Trust (PASSAT), and 2012 International Conference on Social Computing (SocialCom), IEEE, 2012, pp. 31–40.

- [73] A. Passarella, R.I. Dunbar, M. Conti, F. Pezzoni, Ego network models for future internet social networking environments, Comput. Commun. 35 (18) (2012) 2201–2217, doi:10.1016/j.comcom.2012.08.003. URL: http://www. sciencedirect.com/science/article/pii/S0140366412002642
- [74] A. Epasto, S. Lattanzi, V. Mirrokni, I.O. Sebe, A. Taei, S. Verma, Ego-net community mining applied to friend suggestion, Proc. VLDB Endow. 9 (4) (2015) 324–335.
- [75] R.N. Lichtenwalter, J.T. Lussier, N.V. Chawla, New perspectives and methods in link prediction, in: Proceedings of the Sixteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2010, pp. 243–252.
- [76] M. Fire, L. Tenenboim, O. Lesser, R. Puzis, L. Rokach, Y. Elovici, Link prediction in social networks using computationally efficient topological features, in: Proceedings of the 2011 IEEE Third International Conference on Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third International Conference on Social Computing (SocialCom), IEEE, 2011, pp. 73–80.
- [77] L. Backstrom, J. Leskovec, Supervised random walks: predicting and recommending links in social networks, in: Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, ACM, 2011, pp. 635–644.
- [78] P. Victor, C. Cornelis, M. De Cock, A. Teredesai, A comparative analysis of trust-enhanced recommenders for controversial items., in: Proceedings of the 2009 International AAAI Conference on Web and Social Media (ICWSM), 2009.
- [79] J. He, W.W. Chu, A social network-based recommender system (SNRS), in: Data Mining for Social Network Data, Springer, 2010, pp. 47–74.
- [80] J. Golbeck, Generating Predictive Movie Recommendations from Trust in Social Networks, in: K. Stølen, W.H. Winsborough, F. Martinelli, F. Massacci (Eds.), Trust Management: 4th International Conference, iTrust 2006, Pisa, Italy, May 16-19, 2006. Proceedings, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 93–104.
- [81] P. Massa, P. Avesani, Trust metrics on controversial users: balancing between tyranny of the majority, Int. J. Semant. Web Inf. Syst. (IJSWIS) 3 (1) (2007) 39–64.
- [82] M. Jamali, M. Ester, TrustWalker: a random walk model for combining trust-based and item-based recommendation, in: Proceedings of the Fifteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2009, pp. 397–406.
- [83] H. Ma, H. Yang, M.R. Lyu, I. King, SoRec: social recommendation using probabilistic matrix factorization, in: Proceedings of the Seventeenth ACM Conference on Information and Knowledge Management, ACM, 2008, pp. 931– 940.
- [84] M. Jamali, M. Ester, A matrix factorization technique with trust propagation for recommendation in social networks, in: Proceedings of the Fourth ACM Conference on Recommender Systems, ACM, 2010, pp. 135–142.
- [85] H. Ma, D. Zhou, C. Liu, M.R. Lyu, I. King, Recommender systems with social regularization, in: Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, ACM, 2011, pp. 287–296.
- [86] D.D. Lee, H.S. Seung, Learning the parts of objects by non-negative matrix factorization, Nature 401 (6755) (1999) 788.
- [87] R. Salakhutdinov, A. Mnih, Bayesian probabilistic matrix factorization using Markov chain monte carlo, in: Proceedings of the Twenty-fifth International Conference on Machine Learning, ACM, 2008, pp. 880–887.
- [88] H. Ma, I. King, M.R. Lyu, Learning to recommend with social trust ensemble, in: Proceedings of the Thirty-second International ACM SIGIR Conference on Research and Development in Information Retrieval, ACM, 2009, pp. 203–210.
- [89] J.S. Breese, D. Heckerman, C. Kadie, Empirical analysis of predictive algorithms for collaborative filtering, in: Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann Publishers Inc., 1998, pp. 43–52.
- [90] Z. Sun, L. Han, W. Huang, X. Wang, X. Zeng, M. Wang, H. Yan, Recommender systems based on social networks, J. Syst. Softw. 99 (2015) 109–119.
- [91] W. Reafee, N. Salim, A. Khan, The power of implicit social relation in rating prediction of social recommender systems, PloS One 11 (5) (2016) e0154848.
- [92] S. James, C.-H. Michael, The folksonomy tag cloud: when is it useful? J. Inf. Sci. 34 (1) (2008).
- [93] A. Hotho, R. Jäschke, C. Schmitz, G. Stumme, Information retrieval in folksonomies: Search and ranking, in: Proceedings of the 2006 European Semantic Web Conference (ESWC), vol. 4011, Springer, 2006, pp. 411–426.
- [94] Z.-K. Zhang, T. Zhou, Y.-C. Zhang, Tag-aware recommender systems: a state-of-the-art survey, J. Comput. Sci. Technol. 26 (5) (2011) 767–777.
- [95] A.K. Milicevic, A. Nanopoulos, M. Ivanovic, Social tagging in recommender systems: a survey of the state-of-the-art and possible extensions, Artif. Intell. Rev. 33 (3) (2010) 187–209.
- [96] H. Kim, H.-J. Kim, A framework for tag-aware recommender systems, Expert Syst. Appl. 41 (8) (2014) 4000–4009.
- [97] Z. Xu, Y. Fu, J. Mao, D. Su, Towards the semantic web: collaborative tag suggestions, in: Proceedings of the 2006 Collaborative Web Tagging Workshop at WWW2006, Edinburgh, Scotland, 2006.
- [98] B. Sigurbjörnsson, R. van Zwol, Flickr tag recommendation based on collective knowledge, in: Proceedings of the Seventeenth International Conference on World Wide Web, ACM, 2008, pp. 327–336.
- [99] C.S. Firan, W. Nejdl, R. Paiu, The benefit of using tag-based profiles, in: Proceedings of the 2007 Web Conference on Latin American (LA-WEB), IEEE, 2007, pp. 32–41.

- [100] I. Guy, N. Zwerdling, I. Ronen, D. Carmel, E. Uziel, Social media recommendation based on people and tags, in: Proceedings of the Thirty-third International ACM SIGIR Conference on Research and Development in Information Retrieval, ACM, 2010, pp. 194–201.
- [101] J. Peng, D.D. Zeng, H. Zhao, F.-y. Wang, Collaborative filtering in social tagging systems based on joint item-tag recommendations, in: Proceedings of the Nineteenth ACM International Conference on Information and Knowledge Management, ACM, 2010, pp. 809–818.
- [102] A. Shepitsen, J. Gemmell, B. Mobasher, R. Burke, Personalized recommendation in social tagging systems using hierarchical clustering, in: Proceedings of the 2008 ACM Conference on Recommender Systems, ACM, 2008, pp. 259–266.
- [103] J. Gemmell, A. Shepitsen, B. Mobasher, R. Burke, Personalizing Navigation in Folksonomies Using Hierarchical Tag Clustering, in: I.Y. Song, J. Eder, T.M. Nguyen (Eds.), Data Warehousing and Knowledge Discovery: 10th International Conference, DaWaK 2008 Turin, Italy, September 2-5, 2008 Proceedings, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 196–205.
- [104] H. Liang, Y. Xu, Y. Li, R. Nayak, X. Tao, Connecting Users and Items with Weighted Tags for Personalized Item Recommendations, in: Proceedings of the 21st ACM Conference on Hypertext and Hypermedia, ACM, {New York, NY, USA, 2010, pp. 51–60.
- [105] P. Symeonidis, A. Nanopoulos, Y. Manolopoulos, Tag recommendations based on tensor dimensionality reduction, in: Proceedings of the 2008 ACM Conference on Recommender Systems, ACM, 2008, pp. 43–50.
- [106] L. De Lathauwer, B. De Moor, J. Vandewalle, A multilinear singular value decomposition, SIAM J. Matrix Anal. Appl. 21 (4) (2000) 1253–1278.
- [107] S. Rendle, L. Balby Marinho, A. Nanopoulos, L. Schmidt-Thieme, Learning optimal ranking with tensor factorization for tag recommendation, in: Proceedings of the Fifteenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2009, pp. 727–736.
- [108] D.J. Hand, R.J. Till, A simple generalisation of the area under the ROC curve for multiple class classification problems, Mach. Learn. 45 (2) (2001) 171–186.
- [109] S. Brin, L. Page, Reprint of: the anatomy of a large-scale hypertextual web search engine, Comput. Netw. 56 (18) (2012) 3825–3833.
- [110] Z.-K. Zhang, T. Zhou, Y.-C. Zhang, Personalized recommendation via integrated diffusion on user-item-tag tripartite graphs, Phys. A Stat. Mech. Appl. 389 (1) (2010) 179–186.
- [111] V. Arnaboldi, M.G. Campana, F. Delmastro, E. Pagani, Pliers: a popularity-based recommender system for content dissemination in online social networks, in: Proceedings of the Thirty-first Annual ACM Symposium on Applied Computing, ACM, 2016, pp. 671–673.
- [112] V. Arnaboldi, M.G. Campana, F. Delmastro, E. Pagani, A personalized recommender system for pervasive social networks, Pervasive Mob. Comput. 36 (2017) 3–24.
- [113] Y. Zuo, J. Zeng, M. Gong, L. Jiao, Tag-aware recommender systems based on deep neural networks, Neurocomputing 204 (2016) 51–60.
- [114] G.E. Hinton, R.R. Salakhutdinov, Reducing the dimensionality of data with neural networks, Science 313 (5786) (2006) 504–507.
- [115] Z. Xu, T. Lukasiewicz, C. Chen, Y. Miao, X. Meng, Tag-aware personalized recommendation using a hybrid deep model, in: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17, 2017, pp. 3196–3202, doi:10.24963/ijcai.2017/446.
- [116] D. Quercia, N. Lathia, F. Calabrese, G. Di Lorenzo, J. Crowcroft, Recommending social events from mobile phone location data, in: Proceedings of the Tenth IEEE International Conference on Data Mining (ICDM), IEEE, 2010, pp. 971–976.
- [117] J. Bao, Y. Zheng, D. Wilkie, M. Mokbel, Recommendations in location-based social networks: a survey, GeoInformatica 19 (3) (2015) 525–565.
- [118] K. Kodama, Y. Iijima, X. Guo, Y. Ishikawa, Skyline queries based on user locations and preferences for making location-based recommendations, in: Proceedings of the 2009 International Workshop on Location Based Social Networks, ACM, 2009, pp. 9–16.
- [119] M.-H. Park, J.-H. Hong, S.-B. Cho, Location-based recommendation system using Bayesian user's preference model in mobile devices, in: Proceedings of the 2007 International Conference on Ubiquitous Intelligence and Computing, Springer, 2007, pp. 1130–1139.
- [120] L. Ramaswamy, P. Deepak, R. Polavarapu, K. Gunasekera, D. Garg, K. Visweswariah, S. Kalyanaraman, Caesar: a context-aware, social recommender system for low-end mobile devices, in: Proceedings of the Tenth International Conference on Mobile Data Management: Systems, Services and Middleware, MDM'09, IEEE, 2009, pp. 338–347.
- [121] W.R. Tobler, A computer movie simulating urban growth in the detroit region, Econ. Geogr. 46 (sup1) (1970) 234–240.
- [122] T. Horozov, N. Narasimhan, V. Vasudevan, Using location for personalized poi recommendations in mobile environments, in: Proceedings of the 2006 International Symposium on Applications and the Internet, SAINT 2006, IEEE, 2006, pp. 6–pp.
- [123] C.-Y. Chow, J. Bao, M.F. Mokbel, Towards location-based social networking services, in: Proceedings of the Second ACM SIGSPATIAL International Workshop on Location Based Social Networks, ACM, 2010, pp. 31–38.
- [124] Y. Shi, P. Serdyukov, A. Hanjalic, M. Larson, Personalized landmark recommendation based on geotags from photo sharing sites., Proceedings of the 2011 Eleventh International Conference on Web and Social Media (ICWSM), vol. 11, 2011, pp. 622–625.

- [125] Y. Moshfeghi, B. Piwowarski, J.M. Jose, Handling data sparsity in collaborative filtering using emotion and semantic based features, in: Proceedings of the Thirty-fourth International ACM SIGIR Conference on Research and Development in Information Retrieval, ACM, 2011, pp. 625–634.
- [126] A. Levi, O. Mokryn, C. Diot, N. Taft, Finding a needle in a haystack of reviews: cold start context-based hotel recommender system, in: Proceedings of the Sixth ACM Conference on Recommender Systems, ACM, 2012, pp. 115– 122.
- [127] D. Yang, D. Zhang, Z. Yu, Z. Yu, Fine-grained preference-aware location search leveraging crowdsourced digital footprints from LBSNs, in: Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing, ACM, 2013, pp. 479–488.
- [128] H. Wang, Y. Fu, Q. Wang, H. Yin, C. Du, H. Xiong, A location-sentiment-aware recommender system for both home-town and out-of-town users, in: Proceedings of the Twenty-third ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2017, pp. 1135–1143.
- [129] H. Yin, X. Zhou, B. Cui, H. Wang, K. Zheng, Q.V.H. Nguyen, Adapting to user interest drift for poi recommendation, IEEE Trans. Knowl. Data Eng. 28 (10) (2016) 2566–2581.
- [130] D.M. Blei, A.Y. Ng, M.I. Jordan, Latent Dirichlet allocation, J. Mach. Learn. Res. 3 (Jan) (2003) 993–1022.
- [131] M. Ye, P. Yin, W.-C. Lee, Location recommendation for location-based social networks, in: Proceedings of the Eighteenth SIGSPATIAL International Conference on Advances in Geographic Information Systems, ACM, 2010, pp. 458–461.
- [132] M. Ye, P. Yin, W.-C. Lee, D.-L. Lee, Exploiting geographical influence for collaborative point-of-interest recommendation, in: Proceedings of the Thirtyfourth International ACM SIGIR Conference on Research and Development in Information Retrieval, ACM, 2011, pp. 325–334.
- [133] F. Gasparetti, Personalization and context-awareness in social local search: state-of-the-art and future research challenges, Pervasive Mob. Comput. 38 (2017) 446–473.
- [134] J. Wiese, P.G. Kelley, L.F. Cranor, L. Dabbish, J.I. Hong, J. Zimmerman, Are you close with me? are you nearby?: investigating social groups, closeness, and willingness to share, in: Proceedings of the Thirteenth International Conference on Ubiquitous Computing, ACM, 2011, pp. 197–206.
- [135] D. Liben-Nowell, J. Novak, R. Kumar, P. Raghavan, A. Tomkins, Geographic routing in social networks, Proc. Natl. Acad. Sci. U.S.A. 102 (33) (2005) 11623–11628.
- [136] J. Cranshaw, E. Toch, J. Hong, A. Kittur, N. Sadeh, Bridging the gap between physical location and online social networks, in: Proceedings of the Twelfth ACM International Conference on Ubiquitous Computing, ACM, 2010, pp. 119–128.
- [137] Y. Zheng, L. Zhang, X. Xie, W.-Y. Ma, Mining interesting locations and travel sequences from GPS trajectories, in: Proceedings of the Eighteenth International Conference on World Wide Web, ACM, 2009, pp. 791–800.
- [138] J.J.-C. Ying, W.-C. Lee, M. Ye, C.-Y. Chen, V.S. Tseng, User association analysis of locales on location based social networks, in: Proceedings of the Third ACM SIGSPATIAL International Workshop on Location-Based Social Networks, ACM, 2011, pp. 69–76.
- [139] E. Cho, S.A. Myers, J. Leskovec, Friendship and mobility: user movement in location-based social networks, in: Proceedings of the Seventeenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2011, pp. 1082–1090.
- [140] X. Xiao, Y. Zheng, Q. Luo, X. Xie, Finding similar users using category-based location history, in: Proceedings of the Eighteenth SIGSPATIAL International Conference on Advances in Geographic Information Systems, ACM, 2010, pp. 442–445.
- [141] Y. Zheng, L. Zhang, Z. Ma, X. Xie, W.-Y. Ma, Recommending friends and locations based on individual location history, ACM Trans. Web (TWEB) 5 (1) (2011) 5.
- [142] X. Yu, A. Pan, L.-A. Tang, Z. Li, J. Han, Geo-friends recommendation in GP-S-based cyber-physical social network, in: Proceedings of the 2011 International Conference on Advances in Social Networks Analysis and Mining (ASONAM), IEEE, 2011, pp. 361–368.
- [143] S. Scellato, A. Noulas, C. Mascolo, Exploiting place features in link prediction on location-based social networks, in: Proceedings of the Seventeenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2011, pp. 1046–1054.
- [144] A. Pozdnoukhov, C. Kaiser, Space-time dynamics of topics in streaming text, in: Proceedings of the Third ACM SIGSPATIAL International Workshop on Location-based Social Networks, ACM, 2011, pp. 1–8.
- [145] Z. Yin, L. Cao, J. Han, C. Zhai, T. Huang, Geographical topic discovery and comparison, in: Proceedings of the Twentieth International Conference on World Wide Web, ACM, 2011, pp. 247–256.
 [146] V.W. Zheng, Y. Zheng, X. Xie, Q. Yang, Collaborative location and activity rec-
- [140] v.W. Zheng, Y. Zheng, X. Xie, Q. Yang, Collaborative location and activity recommendations with GPS history data, in: Proceedings of the Nineteenth International Conference on World Wide Web, ACM, 2010, pp. 1029–1038.
- [147] P. Symeonidis, A. Papadimitriou, Y. Manolopoulos, P. Senkul, I. Toroslu, Geosocial recommendations based on incremental tensor reduction and local path traversal, in: Proceedings of the Third ACM SIGSPATIAL International Workshop on Location-Based Social Networks, ACM, 2011, pp. 89–96.
- [148] H. Kawakubo, K. Yanai, GeoVisualRank: a ranking method of geotagged imagesconsidering visual similarity and geo-location proximity, in: Proceedings of the Twentieth International Conference Companion on World Wide Web, ACM, 2011, pp. 69–70.

- [149] A. Silva, B. Martins, Tag recommendation for georeferenced photos, in: Proceedings of the Third ACM SIGSPATIAL International Workshop on Location-Based Social Networks, ACM, 2011, pp. 57–64.
- [150] T. Sandholm, H. Ung, Real-time, location-aware collaborative filtering of web content, in: Proceedings of the 2011 Workshop on Context-Awareness in Retrieval and Recommendation, ACM, 2011, pp. 14–18.
- [151] R. Pálovics, P. Szalai, J. Pap, E. Frigó, L. Köcsis, A.A. Benczúr, Location-aware online learning for top-k recommendation, Pervasive Mob. Comput. 38 (2017) 490–504.
- [152] R. Pálovics, A.A. Benczúr, L. Kocsis, T. Kiss, E. Frigó, Exploiting temporal influence in online recommendation, in: Proceedings of the Eighth ACM Conference on Recommender Systems, ACM, 2014, pp. 273–280.
- [153] M. Harvey, F. Crestani, Long time, no tweets! time-aware personalised hashtag suggestion, in: Proceedings of the 2015 European Conference on Information Retrieval, Springer, 2015, pp. 581–592.
- [154] S. Androutsellis-Theotokis, D. Spinellis, A survey of peer-to-peer content distribution technologies, ACM Comput. Surv. (CSUR) 36 (4) (2004) 335–371.
- [155] P. Han, B. Xie, F. Yang, R. Shen, A scalable P2P recommender system based on distributed collaborative filtering, Expert Syst. Appl. 27 (2) (2004) 203–210.
 [156] B.N. Miller, J.A. Konstan, J. Riedl, PocketLens: toward a personal recommender
- system, ACM Trans. Inf. Syst. (TOIS) 22 (3) (2004) 437–476.
- [157] J.K. Kim, H.K. Kim, Y.H. Cho, A user-oriented contents recommendation system in peer-to-peer architecture, Expert Syst. Appl. 34 (1) (2008) 300– 312.
- [158] Y. Zhao, J. Ye, T. Henderson, The effect of privacy concerns on privacy recommenders, in: Proceedings of the Twenty-first International Conference on Intelligent User Interfaces, ACM, 2016, pp. 218–227.
- [159] S.K.T. Lam, D. Frankowski, J. Riedl, Do You Trust Your Recommendations? An Exploration of Security and Privacy Issues in Recommender Systems, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 14–29. 10.1007/11766155_2
- [160] A. De Spindler, M.C. Norrie, M. Grossniklaus, Collaborative filtering based on opportunistic information sharing in mobile ad-hoc networks, in: Proceedings of the 2007 OTM Confederated International Conferences" On the Move to Meaningful Internet Systems", Springer, 2007, pp. 408–416.
- [161] R. Schifanella, A. Panisson, C. Gena, G. Ruffo, MobHinter: epidemic collaborative filtering and self-organization in mobile ad-hoc networks, in: Proceedings of the 2008 ACM Conference on Recommender Systems, ACM, 2008, pp. 27–34.
- [162] D. Karamshuk, C. Boldrini, M. Conti, A. Passarella, Human mobility models for opportunistic networks, IEEE Commun. Mag. 49 (12) (2011) 157–165.
- [163] L. Del Prete, L. Capra, Differs: a mobile recommender service, in: Proceedings of the 2010 Eleventh International Conference on Mobile Data Management (MDM), IEEE, 2010, pp. 21–26.
- [164] N. Eagle, A.S. Pentland, MIT Reality mining dataset, 2005, (Downloaded from https://crawdad.org/mit/reality/20050701). doi:10.15783/C71S31.
- [165] Y. Zhao, J. Ye, T. Henderson, A robust reputation-based location-privacy recommender system using opportunistic networks, in: Proceedings of the Eighth EAI International Conference on Mobile Computing, Applications and Services, ACM, 2016.
- [166] A.R. Beresford, F. Stajano, Location privacy in pervasive computing, IEEE Pervasive Comput. 2 (1) (2003) 46–55.
- [167] S.K. Lam, J. Riedl, Shilling recommender systems for fun and profit, in: Proceedings of the Thirteenth International Conference on World Wide Web, ACM, 2004, pp. 393–402.
- [168] R. Burke, M.P. O'Mahony, N.J. Hurley, Robust collaborative recommendation, in: Recommender Systems Handbook, Springer, 2015, pp. 961–995.
- [169] I. Parris, T. Henderson, F.B. Abdesslem, St_andrews/locshare dataset, 2011, (Downloaded from https://crawdad.org/st_andrews/locshare/20111012). doi:10.15783/C7WW2F.
- [170] G. Lo Giusto, A.J. Mashhadi, L. Capra, Folksonomy-based reasoning for content dissemination in mobile settings, in: Proceedings of the Fifth ACM Workshop on Challenged Networks, ACM, 2010, pp. 39–46.
- [171] A.J. Mashhadi, S.B. Mokhtar, L. Capra, Habit: leveraging human mobility and social network for efficient content dissemination in delay tolerant networks, in: Proceedings of the 172 IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks & Workshops, WoWMoM 2009, IEEE, 2009, pp. 1–6.
- [172] C. Boldrini, A. Passarella, HCMM: modelling spatial and temporal properties of human mobility driven by users' social relationships, Comput. Commun. 33 (9) (2010) 1056–1074.
- [173] M. Salathé, M. Kazandjieva, J.W. Lee, P. Levis, M.W. Feldman, J.H. Jones, A high-resolution human contact network for infectious disease transmission, Proc. Natl. Acad. Sci. 107 (51) (2010) 22020–22025.
- [174] A. Keränen, J. Ott, T. Kärkkäinen, The one simulator for DTN protocol evaluation, in: Proceedings of the Second International Conference on Simulation Tools and Techniques, ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2009, p. 55.
- [175] F. Ekman, A. Keränen, J. Karvo, J. Ott, Working day movement model, in: Proceedings of the First ACM SIGMOBILE Workshop on Mobility Models, ACM, 2008, pp. 33–40.
- [176] M. Pasinato, C.E. Mello, M.-A. Aufaure, G. Zimbrao, Generating synthetic data for context-aware recommender systems, in: Proceedings of the 2013 BRICS Congress on Computational Intelligence and Eleventh Brazilian Congress on Computational Intelligence (BRICS-CCI & CBIC), IEEE, 2013, pp. 563–567.

- [177] M. del Carmen Rodríguez-Hernández, S. Ilarri, R. Hermoso, R. Trillo-Lado, DataGenCARS: a generator of synthetic data for the evaluation of context-aware recommendation systems, Pervasive Mob. Comput. 38 (2017) 516–541.
- [178] G. Guo, J. Zhang, Z. Sun, N. Yorke-Smith, LibRec: a java library for recommender systems., in: Proceedings of the 2015 Conference on User Modeling, Adaptation, and Personalization (UMAP Workshops), 2015.
- [179] M.D. Ekstrand, M. Ludwig, J.A. Konstan, J.T. Riedl, Rethinking the recommender research ecosystem: reproducibility, openness, and LensKit, in: Proceedings of the Fifth ACM Conference on Recommender Systems, ACM, 2011, pp. 133–140.
- [180] P. Castells, N.J. Hurley, S. Vargas, Novelty and diversity in recommender systems, in: Recommender Systems Handbook, Springer, 2015, pp. 881–918.
- [181] Y. Zheng, B. Mobasher, R. Burke, CarsKit: a java-based context-aware recommendation engine, in: Proceedings of the 2015 IEEE International Conference on Data Mining Workshop (ICDMW), IEEE, 2015, pp. 1668–1671.



Mattia Giovanni Campana is a Ph.D. student in Computer Engineering at the University of Pisa in collaboration with the Institute for Informatics and Telematics of the National Research Council of Italy (CNR). He received a M.S. degree in Computer Science from the University of Milan in 2015. His research activity is related to the study and application of Social Recommender Systems to optimize the data dissemination and content discovery in Opportunistic Networks and Mobile Social Networks.



Franca Delmastro is a researcher of the IIT Institute of the National Research Council of Italy (CNR). Her research interests are in the field of Future Internet and Cyber-Physical systems, working on communication protocols for wireless and opportunistic networks, mobile middleware solutions, recommender systems for online and mobile social networks, context and social-aware mobile applications and personalized m-health solutions. She published several works on international journals, conferences and book chapters in the pervasive and mobile computing area. She participated in the organization of several international conferences. She served as TPC Vice Chair of IEEE PerCom 2017, Workshops co-Chair of IEEE

PerCom 2014, General Co-Chair of Work-in-Progress session of IEEE PerCom 2012, and General Chair of some international workshops. She is currently member of the Editorial Board of International Journal of Communications, Network and System Sciences (IJCNS) and she has been Guest Editor of Special Issues for Elsevier Pervasive and Mobile Computing Journal e Computer Communication Journal.