

Hybrid recommendation methods in complex networksA. Fiasconaro,^{1,*} M. Tumminello,² V. Nicosia,¹ V. Latora,^{1,3} and R. N. Mantegna^{4,5,6}¹*School of Mathematical Sciences, Queen Mary University of London, Mile End Road, London E1 4NS, UK*²*Dipartimento di Scienze Economiche, Aziendali e Statistiche, Università di Palermo, Viale delle Scienze Ed. 13, 90128 Palermo, Italy*³*Dipartimento di Fisica ed Astronomia, Università di Catania and INFN, I-95123 Catania, Italy*⁴*Center for Network Science, Central European University, Nador 9 ut., H-1051, Budapest, Hungary*⁵*Department of Economics, Central European University, Nador 9 ut., H-1051, Budapest, Hungary*⁶*Dipartimento di Fisica e Chimica, Università di Palermo, Viale delle Scienze, Edif. 18, I-90128, Palermo, Italy*

(Received 9 December 2014; revised manuscript received 26 May 2015; published 14 July 2015)

We propose two recommendation methods, based on the appropriate normalization of already existing similarity measures, and on the convex combination of the recommendation scores derived from similarity between users and between objects. We validate the proposed measures on three data sets, and we compare the performance of our methods to other recommendation systems recently proposed in the literature. We show that the proposed similarity measures allow us to attain an improvement of performances of up to 20% with respect to existing nonparametric methods, and that the accuracy of a recommendation can vary widely from one specific bipartite network to another, which suggests that a careful choice of the most suitable method is highly relevant for an effective recommendation on a given system. Finally, we study how an increasing presence of random links in the network affects the recommendation scores, finding that one of the two recommendation algorithms introduced here can systematically outperform the others in noisy data sets.

DOI: [10.1103/PhysRevE.92.012811](https://doi.org/10.1103/PhysRevE.92.012811)

PACS number(s): 89.75.Hc, 89.20.Ff, 05.40.Ca

I. INTRODUCTION

The increasingly important role played by information technology and by the ubiquity and success of web-based retail shops is rapidly transforming our lives and buying patterns and is producing a huge quantity of detailed data sets about customers' preferences and habits. The availability of such data sets has made it possible to study in a quantitative way how people select items in several different scenarios such as how they choose movies to watch, books to read, or food to eat. In most of the cases, the number of different items available on an online retail shop is so large that it is extremely difficult to have a clear idea of the specific products that would better fit the taste of each customer. Hence the necessity to devise intelligent automatic systems to provide useful recommendations, based, for instance, on the knowledge about previous purchases made by users. Given its practical importance, the study of recommendation systems is nowadays a very active research topic, with relevant contributions from different fields including computer science, economics, sociology, complex networks, and engineering [1,2].

The natural framework to represent selection or purchasing patterns is by means of a bipartite graph, namely, a graph consisting of two distinct classes of nodes (respectively associated to users and objects) in which two nodes belonging to different classes are connected by an edge if the corresponding user has chosen or purchased that particular object. Within this framework, a recommendation is no more than the suggestion of (a relatively small) set of objects to which a specific user might be interested and corresponds to a set of new potential edges in the bipartite graph. In many cases, an object is recommended to a user based on her similarity with other

users, so that the definition of appropriate similarity measures is crucial for the development of efficient *personalized recommendation systems*. Various recommendation systems and algorithms have been proposed over the years, such as collaborative filtering (CF) [3–5], methods based on diffusion across the user-object network [6–10], and hybrid (parametric) combinations of different algorithms [9–12]. In most of the cases, the quantification of similarity between two users is based on the number of objects which have been chosen by both users in the past. However, it is also possible to define a similarity between two objects based on the number of users who have chosen them. Not always in the literature have recommendations based on user similarity been properly distinguished from those based on object similarity, and the predictions provided by these two types of recommendation systems have been usually compared while discarding the different nature of the similarity measures involved.

The question of which similarity measure is the most reliable in providing tailored and accurate recommendations is still a matter of open debate, and it is not clear yet how to choose one similarity definition or another for a specific recommendation task. This paper provides a contribution in this direction. In particular, we focus here on the duality user similarity versus object similarity, showing that it is possible to improve the quality of recommendation by making a combined use of the two classes of similarity. We start by proposing two new definitions of similarity based on heuristic arguments, and then we compare the accuracy of recommendations based on these definitions with the accuracy of other methods proposed in the literature. The first measure we propose takes into account the popularity of objects and the heterogeneity of user selection patterns, while the second one is based on the concept of Pearson correlation [13] generalized to the case of binary vectors. We then show that any definition of similarity between users induces a definition of similarity

*A.Fiasconaro@qmul.ac.uk

between objects, and vice versa. This fact actually increases the number of different possible similarity definitions and allows us to consider recommendation methods which combine user similarities with object similarities. We also test the robustness of different similarity measures against the presence of noise, by adding an increasing percentage of random edges to the actual bipartite graph, and we show that the measure based on the generalized Pearson correlation proposed in the paper is able to filter out noise more effectively than most of the other existing similarity measures. In general, a method that might increase the performance of recommendation systems is the so-called dimensionality reduction, which consists of filtering and/or manipulating the data and/or the similarity matrices in order to select the most relevant factors that describe the structure of the data set [14,15]. We present here an example of rank reduction [16] that makes use of the singular value decomposition of the incidence matrix and confirms the robustness of the proposed similarities.

The paper is organized as follows. In Sec. II we review different similarity measures proposed in the literature, we introduce two new definitions of similarity, and we show how to associate a recommendation score to an object starting from a given similarity definition. In Sec. III we provide a brief description of three data sets corresponding to user-object associations in different contexts, and we validate the performance of different recommendation strategies on the corresponding bipartite networks. In Sec. IV we show how a measure of similarity between users can be transformed into an analog measure of similarity between objects. We then investigate recommendation methods which combine the recommendation scores obtained from similarities between users and similarity between objects. In Sec. V we study how the presence of spurious information in the data sets can affect the performance of recommendation, and we show that some recommendation strategies actually perform better than others in noisy data sets. In Sec. VI we analyze the effect of rank reduction of the incidence matrix on the performance of our recommendation methods. Finally in Sec. VII we draw our conclusions.

II. MEASURES OF SIMILARITY

Let us consider a set of N users and M objects, where each user i is associated to a subset of the M objects she has expressed a preference for. This is, for instance, the case of users buying objects from an online retail shop, where each user is associated to each of the items she has bought from that web site. Such systems can be naturally represented by means of bipartite graphs, where users and objects are considered as two distinct classes of nodes. A bipartite graph can be described by an $N \times M$ incidence matrix A whose entry $a_{i,\alpha}$ is equal to 1 if and only if user u_i is associated to object o_α (for instance, because u_i has bought that object) and $a_{i,\alpha} = 0$ otherwise. Notice that in a bipartite graph each edge always connects one user with one object. In the following, Latin subscripts are associated to users, whereas Greek ones are associated to objects. The total number of objects collected by a user u_i is equal to the number of edges incident on the corresponding node i of the graph, i.e., to the degree $k_i = \sum_{\alpha} a_{i,\alpha}$. Similarly, the degree of object o_α , defined as

$k_\alpha = \sum_i a_{i,\alpha}$, is equal to the total number of users that have collected that object.

Within this framework, making a *recommendation* for user u_i corresponds to compiling a list of objects which have not already been chosen (or bought) by user u_i but to which u_i might be interested. In other words, a recommendation is just a proposal of new potential edges of the bipartite graph whose one endpoint is node i . The main hypothesis on which almost all recommendation systems rely is that the set of objects actually collected (or bought) by user u_i represents a sample of her tastes and preferences and can therefore be used to compile a profile of user u_i and to predict which kind of objects u_i might be interested. Consequently, each recommendation systems relies on some measure of similarity. In general, it is possible to define a similarity $s_{i,j}^u$ for the (ordered) pair of users u_i and u_j and also a similarity $s_{\alpha,\beta}^o$ between the pair of objects o_α and o_β , and various different definitions have been proposed in the literature [17].

Similarity measures. A very simple way of quantifying the similarity between user u_i and user u_j is by counting the number of objects $n_{i,j}$ that they have in common:

$$n_{i,j} = \sum_{\alpha=1}^M a_{i,\alpha} a_{j,\alpha}.$$

One of the limitations of $n_{i,j}$ is that it does not take into account the differences in the total number of objects collected by each user. This problem can be somehow alleviated by using the so-called Jaccard similarity [18], defined as the ratio between the number of items collected by both users u_i and u_j , and the sum of the degrees of the two users:

$$s_{i,j}^{u,J} = \frac{\sum_{\alpha=1}^M a_{i,\alpha} a_{j,\alpha}}{(k_i + k_j)} = \frac{n_{i,j}}{(k_i + k_j)}. \quad (1)$$

Another widely used similarity measure is the Hub Promoted Index (HPI), often used in the context of the *collaborative filtering* (CF) approach:

$$s_{i,j}^{u,HPI} = \frac{\sum_{\alpha=1}^M a_{i,\alpha} a_{j,\alpha}}{\min\{k_i, k_j\}}. \quad (2)$$

In Eq. (2), the similarity measure is proportional to the number of objects $n_{i,j}$ users u_i and u_j have in common, and inversely proportional to the smallest of the two degrees, i.e., to $\min\{k_i, k_j\}$. In this way, if user u_i has collected exactly one object, which has also been collected by u_j who instead has degree $k_j \gg 1$, then $s_{i,j}^{u,CF} = 1$, i.e., the similarity between two users is effectively determined by the user with the smallest degree.

The Jaccard and the HPI similarities do not take into account another type of heterogeneity, namely, the fact that not all objects have the same popularity. Intuitively, objects that have been collected by a relative large number of users (in a limiting case, by all users), do not provide useful information for a personalized recommendation, for the simple reason that they are common to too many users, and therefore the fact that one user has collected them does not tell much about her tastes. Hence, it might be a good idea to discount the contribution of an object to the similarity between two users by a function of the degree of the object.

The so called network-based inference (NBI) recommendation method [8] is based on a measure of similarity which takes into account the heterogeneity of users and objects (this recommendation strategy is also called *probabilistic spreading* in a subsequent paper) [10]. In this case, the similarity measure is defined as

$$s_{\alpha,\beta}^{\text{u,NBI}} = \frac{1}{k_\beta} \sum_{l=1}^N \frac{a_{l,\alpha} a_{l,\beta}}{k_l}. \quad (3)$$

This is a similarity between objects, where the contribution of the user u_l which collects the two objects α and β is discounted by the degree of that user k_l , and the whole sum is divided by the degree of one of the two objects, according to the *resource-allocation* procedure defined by the authors in Ref. [8].

It is worth noting that this definition of similarity, like the analogous one $s_{\alpha,\beta}^{\text{o,HeatS}} = \frac{1}{k_\alpha} \sum_{l=1}^N \frac{a_{l,\alpha} a_{l,\beta}}{k_l}$ investigated in Refs. [6,10], is asymmetric, meaning that $s_{\alpha,\beta}^{\text{o,NBI}} \neq s_{\beta,\alpha}^{\text{o,NBI}}$, and $s_{\alpha,\beta}^{\text{o,HeatS}} \neq s_{\beta,\alpha}^{\text{o,HeatS}}$. Though asymmetry is not in general an issue for the recommendation task, it has been shown that better performance can be achieved by using symmetrized versions of these measures [10,12]. Nevertheless, NBI has proved to be a quite reliable recommendation method, and in the following we will consider it as a reference to quantify the effectiveness of the recommendation strategies we propose.

The two new recommendation methods we propose in this paper are based on symmetric similarity measures. Specifically, the first measure we propose is

$$s_{i,j}^{\text{u,MDW}} = \frac{1}{\max\{k_i, k_j\}} \sum_{\alpha=1, k_\alpha > 1}^M \frac{a_{i,\alpha} a_{j,\alpha}}{k_\alpha - 1}. \quad (4)$$

We call it maximum degree weighted (MDW) similarity because the total number of objects collected by both u_i and u_j is weighted by the maximum of the degrees of the two users. Moreover, the contribution of object o_α is weighted by its degree k_α . We note here that some recent studies have investigated the effect of a tunable power-law function of the degree, i.e., of similarity measures in which the contribution due to object o_α is divided by $(k_\alpha)^a$ [6,19,20].

In the following, we will briefly explain the rationale behind Eq. (4). First, the contribution to the similarity measure of each object collected by both users is weighted with the degree of the object. In this way, popular objects will provide smaller contributions to the similarity between users. In particular, the value $k_\alpha - 1$ in the denominator allows us to obtain a maximum contribution to similarity (exactly equal to 1) if and only if $a_{i,\alpha} a_{j,\alpha} = 1$ and $k_\alpha = 2$. This takes into account the very special case in which u_i and u_j are the only two users who have collected a certain object. Second, the similarity measure is divided by the maximum of the degrees of the two users. As we explained above, this choice allows us to properly take into account the existing heterogeneity in the number of selection made by each user. For instance, if we consider the similarity defined in Eq. (2) and we assume that users u_i and u_j have degree $k_i = 1$ and $k_j = 100$ and have exactly one object in common, then the contribution of that object to the similarity between the two users would be equal to 1. However, the contribution of the only object in common between u_i and u_j is equal to 1 also when $k_i = 1$ and $k_j = 2$,

despite that one would argue that in the latter case the two users are more similar than in the former.

By dividing for the maximum of the degrees of the two users, the similarity measure given in Eq. (4) assigns a higher value of similarity to the two users in the latter case (when $k_i = 1$ and $k_j = 2$ we get $s_{i,j}^{\text{u}} = 1/2$) than in the former case (i.e., when $k_i = 1$ and $k_j = 100$, for which we obtain $s_{i,j}^{\text{u}} = 1/100$).

A second similarity measure we propose here is based on the Pearson correlation coefficient between binary vectors and is defined as follows [21]:

$$s_{i,j}^{\text{u,BP}} = \frac{n_{i,j} - k_i k_j / M}{\sqrt{k_i(1 - k_i/M)k_j(1 - k_j/M)}}. \quad (5)$$

This measure, which is denoted in the following as binary Pearson (BP) similarity, is based on the fact that the i th row \mathbf{a}_i of the incidence matrix A represents the profile vector of user u_i , i.e., the set of objects selected by the user. If we have two users, u_i and u_j , who have collected k_α and k_β objects in total, respectively, and we consider the corresponding profile vectors \mathbf{a}_i and \mathbf{a}_j , then we have

$$\langle a_i \rangle = \frac{1}{M} \sum_{\alpha=1}^M a_{i,\alpha} = \frac{k_i}{M}, \quad (6)$$

$$\langle a_i^2 \rangle = \frac{1}{M} \sum_{\alpha=1}^M a_{i,\alpha}^2 = \frac{1}{M} \sum_{\alpha=1}^M a_{i,\alpha} = \frac{k_i}{M}, \quad (7)$$

$$\langle a_i a_j \rangle = \frac{1}{M} \sum_{\alpha=1}^M a_{i,\alpha} a_{j,\alpha} = \frac{n_{i,j}}{M}, \quad (8)$$

where $n_{i,j}$ is the number of objects collected by both u_i and u_j . Therefore the Pearson's correlation coefficient between the two vectors is

$$\begin{aligned} s_{i,j}^{\text{u,BP}} &= \frac{\langle a_i a_j \rangle - \langle a_i \rangle \langle a_j \rangle}{\sqrt{(\langle a_i^2 \rangle - \langle a_i \rangle^2)(\langle a_j^2 \rangle - \langle a_j \rangle^2)}} \\ &= \frac{\frac{n_{i,j}}{M} - \frac{k_i k_j}{M^2}}{\sqrt{\left(\frac{k_i}{M} - \frac{k_i^2}{M^2}\right)\left(\frac{k_j}{M} - \frac{k_j^2}{M^2}\right)}} \\ &= \frac{n_{i,j} - k_i k_j / M}{\sqrt{k_i(1 - k_i/M)k_j(1 - k_j/M)}}, \end{aligned} \quad (9)$$

which is identical to Eq. (5).

The Pearson's correlation coefficient has been used in the literature of recommendation methods in order to quantify similarities based on item ratings [22,23]. Quite differently, the Pearson correlation coefficient is applied here to binary data, obtained by assigning 1 to objects that have been selected by the users and 0 otherwise. The so-called BP similarity presents some remarkable properties. First, it is invariant with respect to the scaling of the system, i.e., with respect to a transformation $k_i \rightarrow q k_i$, $k_j \rightarrow q k_j$, $n_{i,j} \rightarrow q n_{i,j}$, and $M \rightarrow q M$, where q is a positive integer. Furthermore, $s_{i,j}^{\text{u,BP}}$ can be interpreted in terms of the hypergeometric distribution $H(X, k_i, k_j, M)$. Indeed, the mean value of $H(X, k_i, k_j, M)$ is $m_{i,j} = k_i k_j / M$, and the variance is $\sigma_{i,j}^2 = \frac{1}{M-1} k_i(1 - k_i/M)k_j(1 - k_j/M)$. Therefore $s_{i,j}^{\text{u,BP}}$ is proportional to the standard score $z_{i,j}^H$

associated with observation $n_{i,j}$ according to the hypergeometric distribution [24]:

$$s_{i,j}^{u,BP} = \frac{1}{\sqrt{M-1}} \frac{n_{i,j} - m_{i,j}}{\sigma_{i,j}} = \frac{1}{\sqrt{M-1}} z_{i,j}^H.$$

This equation reveals that $s_{i,j}^{u,BP}$ is conceptually different from all the other similarity measures introduced above. Indeed, according to $s_{i,j}^{u,BP}$, the similarity between two users does not depend only on the number of objects selected by both users, $n_{i,j}$, but it depends on the difference between $n_{i,j}$ and the number of shared objects that is expected under the hypothesis that the two users have picked the objects at random. Therefore $s_{i,j}^{u,BP}$ can also take negative values, and this fact influences the way in which a personalized recommendation value is obtained, as discussed in the reminder of this section.

Constructing recommendation lists. Once we have assigned a similarity value to each possible pair of users in the system, using a certain similarity measure, we need an algorithm to construct a recommendation list, i.e., a list of suggested objects which have not been yet collected by a certain user u_i . The simplest way of constructing a recommendation list is the *global ranking method* (GRM). It consists in creating a user recommendation list by considering all the objects not collected by user u_i in decreasing order of their degree. This method is not personalized, except for the fact that objects already collected by that user are excluded from the corresponding list.

A more effective and widely used basic procedure is *collaborative filtering* (CF), which is based on the similarity measure HPI given in Eq. (2). The similarity scores between user u_i and all the other users in the system are used to construct a personalized recommendation value $v_{i,\alpha}^u$, that is, an estimation of how much user u_i might be interested in object α :

$$v_{i,\alpha}^u = \frac{\sum_{l=1, l \neq i}^N s_{i,l}^u a_{l,\alpha}}{\sum_{l=1, l \neq i}^N |s_{i,l}^u|}. \quad (10)$$

The presence of the absolute value at the denominator is not necessary for many of the similarities presented above, being their values always positive. The only exception is the BP similarity, which may take both positive and negative values, thus requiring a proper normalization to avoid possible divergences. In the NBI recommendation method, the recommendation value $f_{i,\alpha}^o$ is defined in a quite different way:

$$f_{i,\alpha}^o = \sum_{\beta=1}^M s_{\alpha,\beta}^{o,NBI} a_{i,\beta}. \quad (11)$$

In fact the computation of the recommendation value includes self-similar terms ($s_{\alpha,\alpha}^o$) which are not taken into account in Eq. (10) and do not include any additional normalization factor.

III. DATA SETS AND VALIDATION

We considered three classical data sets of user-object associations, namely, the MovieLens database, where N users have rated M movies, the Jester Jokes database, where we find records of users who have rated jokes, and the Fine Foods database, containing Amazon reviews of fine foods. In all these

TABLE I. Summary statistics of the three databases used.

	N_{links}	$N \times M$ (used)	Rating range	Thre- shold	N_{links} Filtered
MovieLens	100 K	$943 \times 1,681$	[1, 5]	3	90 K
Jester Jokes	141 K	2000×100	[-10, 10]	0	57 K
Fine Foods	95 K	2000×3317	[1, 5]	3	83 K

databases users have rated items with an ordinal attribute [25]. In our study we will perform recommendation procedures on the incidence matrix of the corresponding bipartite network. For this reason, similarly to what is adopted in other studies (see Ref. [8]), we assume that a user has collected an object if and only if she has rated the object with a score higher or equal to a certain threshold. As a consequence, the databases are filtered with the most appreciated objects, without any further distinction among the highest ratings. In Table I we report information about the size of each data set, the rating range, and the values of the threshold used to filter the ratings and construct an unweighted bipartite network.

Distributions of similarity values. As a preliminary investigation, we evaluated the heterogeneity of the degree of users and objects in the three databases. In Fig. 1 we show the degree distributions for the three databases used. With the only exception of the degree distribution of jokes (objects) in the Jester Jokes database, all the distributions exhibit relatively broad tails. This suggests that similarity measures which properly take into account degree heterogeneities should indeed provide better recommendations.

As a matter of fact, different similarity measures produce different distribution of similarity scores. In Fig. 2 and Fig. 3 we show the probability density functions of the MDW and BP similarities measures on the three investigated databases. The top panels (a), (b), and (c) of each figure report the distribution of user similarities, whereas the bottom panels (d), (e), and (f) correspond to object similarity. It is worth noting that the profile of the probability density functions is

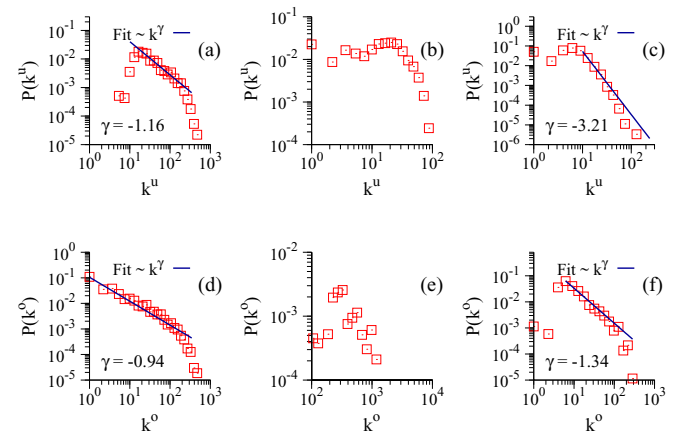


FIG. 1. (Color online) Degree distributions of users [top panels (a), (b), (c)] and objects [bottom panels (d), (e), (f)] of the three data sets, respectively: MovieLens [left panels (a) and (e)], Jester Jokes [middle panels (b) and (f)], and Fine Food [right panels (c) and (f)].

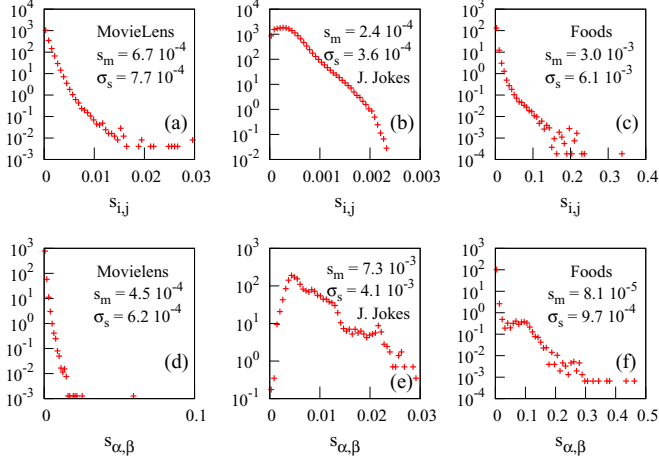


FIG. 2. (Color online) Distribution of the MDW similarity values for the three databases. The top panels show the probability density functions of users' similarity, whereas the bottom panels show the probability density functions of object similarity. Left, middle, and right columns refer to MovieLens [(a), (d)], Jester Jokes [(b), (e)], and Fine Foods [(c), (f)] databases, respectively.

strongly dependent on the similarity measure adopted and is qualitatively different in the three data sets.

Validation. In order to compare the performance of the proposed similarity measures with those of other existing similarity definitions, we split each data set into two sections. Starting from the incidence matrix $A = \{a_{i,\alpha}\}$ representing all the user-object associations in a data set, we considered a subgraph $T = \{t_{i,\alpha}\}$ to be used to compute the similarity scores and recommendation lists for all the users (the so-called *training set*), while the remaining subgraph $W = \{a_{i,\alpha}\} \setminus \{t_{i,\alpha}\} = \{w_{i,\alpha}\}$ was used for *validation*. The recommendation lists obtained from the training sets are compared with the object selections included in the validation set, in order to

check whether users have actually collected objects which are ranked high in their recommendation lists. In the following we report the results corresponding to training sets containing 90% of all the edges of each data set, chosen at random, while the validation sets consist of the remaining 10% of edges. Qualitatively similar results were obtained for different compositions of the training and validation sets.

A basic measure to quantify the performance of a recommendation method is the *rank quality index* r , which is computed as the average quality of recommendation over all the users of the data set. For each user u_i we define the quality of the recommendation r_i provided to u_i as the average of the ratio

$$r_{i,\alpha} = w_{i,\alpha} \frac{L_{i,\alpha}}{M - k_i} \quad (12)$$

computed over all the objects in the recommendation list of user u_i which have actually been selected by u_i in the validation set. Here $L_{i,\alpha}$ is the position of o_α in the recommendation list of u_i , where $L_{i,\alpha} = 1, 2, \dots$, if the object o_α is ranked first, second, \dots etc. in the recommendation list of u_i . Consequently, better recommendations are associated to smaller values of $r_i = \sum_\alpha r_{i,\alpha}$. As we anticipated above, the rank quality index r is the average of $r_{i,\alpha}$ over all the users in the data set:

$$r = \langle r_{i,\alpha} \rangle = \frac{\sum_{i=1}^N \sum_{\alpha=1}^M r_{i,\alpha}}{\sum_{i=1}^N \sum_{\alpha=1}^M w_{i,\alpha}}. \quad (13)$$

Another validation measure testing the accuracy of the predictions is the *hitting rate*, $\text{hit}(L)$, i.e., for all the users, the ratio between the number of collected objects included in the recommendation list of length L , and the number of objects effectively collected *up to* the possible maximum value $N \cdot L$. While $\text{hit}(L)$ represents a percentage of the objects a user has collected among recommended ones, the r value takes into account the position of collected objects in the recommendation list. According to these definitions, a good recommendation method should minimize the value of r and maximize the value of $\text{hit}(L)$.

In the following analysis we set $L = 20$. This value represents a reasonable choice of the length of a recommendation list that can be proposed to a user. Moreover, it has been already used in other works and can be considered a valid reference to compare the performances of different recommendation methods.

For each data set, we considered $N_e = 20$ independent realizations of the training set T , obtained by selecting uniformly at random 90% of the edges in the data set, we constructed the recommendation list induced by each similarity measure, and we computed the value of the rank quality index r and of the hitting rate $\text{hit}(20)$. In the following we report the average values of r and $\text{hit}(20)$ and their associated statistical errors (the standard deviations of the means), respectively, denoted by $\langle r \rangle$, $\langle \text{hit}(20) \rangle$, $\sigma_{\langle r \rangle}$, and $\sigma_{\langle \text{hit}(20) \rangle}$. The mean values $\langle r \rangle$ and $\langle \text{hit}(20) \rangle$ obtained over the $N_e = 20$ different realizations are shown in Table II.

By analyzing the outcomes summarized in Table II we see that the best results are obtained by different methods in different databases. Moreover the two indicators $\langle r \rangle$ and $\langle \text{hit}(20) \rangle$ always single out a different method as the best one.




FIG. 3. (Color online) Distribution of the BP similarity values for the three databases. The top panels show the probability density functions of users' similarity, whereas the bottom panels show the probability density functions of object similarity. Left, middle, and right columns refer to MovieLens [(a), (d)], Jester Jokes [(b), (e)], and Fine Foods [(c), (f)] databases, respectively.

TABLE II. Average rank quality index $\langle r \rangle$ and hitting rate $\langle \text{hit}(20) \rangle$ for the different recommendation methods on each of the three data sets. The mean is computed on $N_e = 20$ different realizations. The standard deviation of the mean values is given in parenthesis. For each database we highlight in boldface the best result. The column Δ reports the difference of the measure with respect to NBI in percentage.

MovieLens	$\langle r \rangle$	$\sigma_{\langle r \rangle}$	$\Delta(\%)$	$\langle \text{hit}(20) \rangle$	$\sigma_{\langle \text{hit}(20) \rangle}$	$\Delta(\%)$
GRM	0.13821 (0.00038)		31.4	0.1928	(0.0041)	-29.4
CF	0.11882 (0.00037)		13.0	0.2364	(0.0010)	-13.4
NBI	0.10514 (0.00028)		0.0	0.2732	(0.0010)	0.0
MDW	0.10563 (0.00022)		0.47	0.2766	(0.0010)	1.2
BP	0.10728 (0.00032)		2.0	0.2708	(0.0009)	-0.9
J	0.11442 (0.00035)		8.8	0.2568	(0.0010)	-6.0
<i>Jester Jokes</i>						
GRM	0.30332 (0.00049)		6.7	0.6160	(0.0006)	-9.1
CF	0.28718 (0.00051)		1.0	0.6712	(0.0008)	-0.9
NBI	0.28422 (0.00045)		0.0	0.6775	(0.0012)	0.0
MDW	0.28087 (0.00037)		-1.2	0.6806	(0.0010)	0.5
BP	0.23795 (0.00052)		-16.3	0.6653	(0.0012)	-1.8
J	0.28549 (0.00049)		0.45	0.6716	(0.0014)	-0.9
<i>Fine Foods</i>						
GRM	0.22263 (0.00073)		1700	0.0891	(0.0004)	-87.9
CF	0.01458 (0.00021)		18.5	0.7000	(0.0014)	-5.4
NBI	0.01230 (0.00012)		0.0	0.7402	(0.0013)	0.0
MDW	0.01304 (0.00017)		6.0	0.7293	(0.0005)	-1.4
BP	0.01173 (0.00010)		-4.6	0.5777	(0.0009)	-21.9
J	0.01534 (0.00013)		24.7	0.6990	(0.0009)	-5.6

However, an overall analysis shows that the best recommendation methods are NBI [the best method according to $\langle r \rangle$ in the MovieLens database and the best method according to $\langle \text{hit}(20) \rangle$ in the Fine Foods database], MDW [the best method according to $\langle \text{hit}(20) \rangle$ in the MovieLens and Fine Foods databases], and BP (the best method according to $\langle r \rangle$ in the Jester Jokes and in the Fine Foods databases). They clearly overcome the results obtained by GRM, CF, and Jaccard (J).

IV. HYBRID OBJECT-USER METHODS

The most important difference between the recommendation methods compared in Table II is that while NBI is based on a definition of similarity among objects, all the other methods make use of similarity measures defined between users.

In general, it is possible to define a transformation rule to obtain a similarity score between users starting from a similarity between objects, and vice versa. In fact, the similarity between objects $s_{i,j}^o$ can be obtained from the similarity between users by appropriately swapping Latin indexes with Greek ones, and quantities defined for users with the analogous ones defined for objects:

$$s_{i,j}^u \leftrightarrow s_{\alpha,\beta}^o, \quad i,j \leftrightarrow \alpha,\beta, \quad N \leftrightarrow M. \quad (14)$$

The transformation rule is valid in both directions from user to objects and from objects to users.

We propose to define new recommendation scores by using the dual similarity measures obtained with the above defined transformation. For example, the recommendation

value, which is the dual of Eq. (10) and is valid for objects instead of users, is obtained as

$$v_{i,\alpha}^o = \frac{\sum_{\beta=1, \beta \neq \alpha}^M s_{\alpha,\beta}^o \cdot a_{i,\beta}}{\sum_{\beta=1, \beta \neq \alpha}^M |s_{\alpha,\beta}^o|}, \quad (15)$$

whereas the dual recommendation score of the NBI algorithm is

$$f_{i,\alpha}^u = \sum_{l=1}^N s_{i,l}^{u,\text{NBI}} \cdot a_{l,\alpha}. \quad (16)$$

It is interesting to note that according to the definition of the NBI we have

$$f_{i,\alpha}^o = f_{i,\alpha}^u = f_{i,\alpha}. \quad (17)$$

This relation can be verified by replacing $s_{\alpha,\beta}^{o,\text{NBI}}$ [Eq. (3)] with $s_{i,l}^{u,\text{NBI}} = \frac{1}{k_l} \sum_{\beta=1}^M \frac{a_{i,\beta} a_{l,\beta}}{k_\beta}$ into Eqs. (11) and (16), respectively. Hence, NBI is invariant under the transformation rules of Eq. (14). It is interesting to investigate how the duality user-object similarity affects the quality of recommendation. To this aim, we propose to define a recommendation value $v_{i,\alpha}(\lambda)$ which is the result of the convex combination of the two recommendation values $v_{i,\alpha}^u$ and $v_{i,\alpha}^o$ obtained from the similarity between users and between objects, respectively. In formula:

$$v_{i,\alpha}(\lambda) = (1 - \lambda)v_{i,\alpha}^u + \lambda v_{i,\alpha}^o, \quad (18)$$

where the relative weight of the user and object recommendation values is controlled by the parameter $\lambda \in [0, 1]$, so that when $\lambda = 0$ we recover the recommendation score induced by the similarity between users, while for $\lambda = 1$ we have the recommendation score corresponding to the similarity between objects. Our hypothesis, which is validated in the following, is that better recommendations can be obtained by appropriately tuning the value of λ .

The mean values $\langle r \rangle$ for different recommendation methods are reported in Fig. 4, where the three panels show the results obtained in the three data sets. It is worth noting that the NBI algorithm is independent of λ . In fact, by using Eq. (17) one verifies that

$$f_{i,\alpha}(\lambda) = (1 - \lambda)f_{i,\alpha}^u + \lambda f_{i,\alpha}^o = f_{i,\alpha}. \quad (19)$$

In Fig. 4 we notice that the CF recommendation method performs poorly for almost all the values of λ , in all the three data sets. In the case of MovieLens, three recommendation methods (MDW, NBI, and BP) perform in a similar way when only the user similarity measure is taken into account ($\lambda = 0$) as we already noticed in the results summarized in Table II. On the other hand, for $\lambda = 1$, i.e., when only the object similarity measure is taken into account, the MDW method performs better than the others. In the case of the Fine Foods data set, the BP similarity performs slightly better than the others for $\lambda = 0$ and for a relatively large range of λ values. When $\lambda = 1$, the recommendation with the MDW measure performs slightly better. Finally, in the Jester Jokes data set the BP similarity clearly outperforms all the others when $\lambda = 0$, while for $\lambda = 1$ all the methods provide similar results, with the only exception of the CF recommendation, whose performance is much worse.

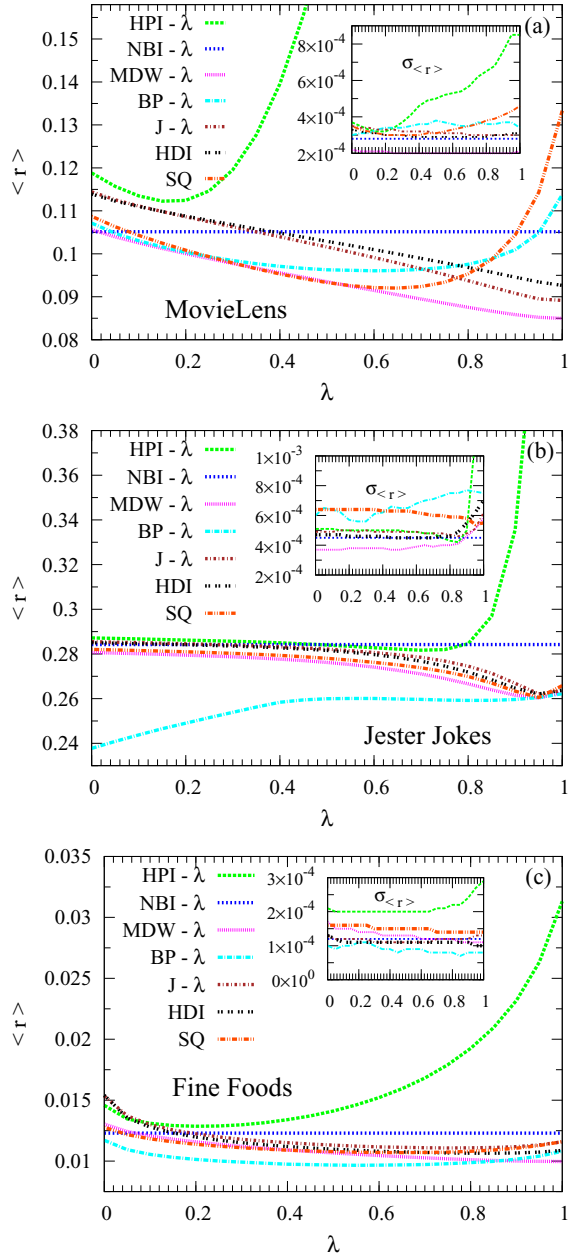


FIG. 4. (Color online) Average rank quality index $\langle r \rangle$ for the different recommendation methods as a function of the user-object parameter λ . The insets report the corresponding standard errors of the mean.

The variety of the profiles observed in Fig. 4 suggests that the performance of a recommendation method depends both on the specific database and on the specific linear combination of user and object recommendation values adopted. Quite often, the best recommendation is not the one corresponding to $\lambda = 0$ or $\lambda = 1$. Some methods perform better at the user limit ($\lambda = 0$), others at the object limit ($\lambda = 1$), some of them for an intermediate value of λ . Moreover, the specific shape of $\langle r \rangle$ as a function of λ actually depends on the database. We would like to stress two interesting aspects of these results. First, some methods exhibit a convex profile of $\langle r \rangle$ as a function of λ , where the minimum indicates the best linear combination of user and

object recommendation values. Second, the variability of the values of $\langle r \rangle$ obtained by different recommendation systems is much higher for $\lambda = 1$ than for $\lambda = 0$.

We have also compared BP and MDW similarities to two additional similarity measures, the Hub Depressed Index (HDI), $s_{i,j}^{u,HDI} = \frac{\sum_{\alpha=1}^M a_{i,\alpha} a_{j,\alpha}}{\max\{k_i, k_j\}}$ [17], and the SQ measure $s_{i,j}^{u,SQ} = \frac{1}{\sqrt{k_i k_j}} \sum_{\alpha=1, k_\alpha > 0}^M \frac{a_{i,\alpha} a_{j,\alpha}}{k_\alpha}$ [19].

These results are also reported in Fig. 4 for the three data sets and can be summarized as follows. The performances obtained by using the HDI show similar behavior as the Jaccard measure, just a little worse, but much better than HPI, and so validating the use of the maximum value of the user degree in the factor of the similarity with respect to the HPI. As in HPI, HDI similarity does not take into account any object weight inside the sum. In the case of the SQ similarity weighted in the objects [19], the performance depends on λ in a way which is comparable with Jaccard measure in both Fine Foods and Jester Jokes databases, while a pronounced nonmonotonic behavior appears in MovieLens database, with a low performance at the object limit ($\lambda = 1$). By comparing the shape of these two similarities with MDW, we conclude that the measure we propose appears somehow as a combination of the two discussed approaches: weight with the object degree inside the sum, and weight according to the maximum user degree outside the sum.

V. IMPACT OF RANDOMNESS

In this section we analyze the robustness of recommendation systems against the presence of different sources of noise in the data sets. Randomness in the data can be originated by different factors, such as errors in object selection, wrong transcription of the data, and fluctuations in the data of a specific kind of objects, for example, due to high heterogeneity of users. We consider here three different ways to include noise in the bipartite networks under study. In the first case we add a certain amount of random edges to the bipartite graph, mimicking erroneously reported user selections. In the second case we rewire a given percentage of the edges of the bipartite network by maintaining the degree of users unaltered (while the degree distribution of object is in general modified). Finally, in the third case we rewire a fraction of the edges of the graph by maintaining unaltered both the user and object degree distribution.

For the sake of simplicity, we show the results obtained for the three randomizing methods only for the MovieLens database. In Fig. 5 we show the average rank quality index for the different methods with $\lambda = 0$ as a function of the percentage of edges randomly added or rewired. As expected, $\langle r \rangle$ is an increasing function of the percentage of noise, signaling a degradation of the recommendation performance. However, the actual profile of $\langle r \rangle$ depends on the specific recommendation method used. In fact, several curves crosses at different values of the induced randomness. This is clearly observed for the first and second kinds of randomization.

We performed the same analysis also on the Jester Jokes and Fine Foods databases, and we report in Fig. 6 the results corresponding to the first type of randomization (addition of a




FIG. 5. (Color online) Mean validation values as a function of noise in the MovieLens database. The three panels correspond, respectively, to the addition of edges at random (a), to edge rewiring (b) and (c) maintaining unaltered the only user degree sequence (b), or maintaining both the user and object degree sequences (c).

fraction of random edges). The results show a prominent role of the BP similarity measure, which seems the most robust in dealing with noisy data sets.

Our findings suggest that the BP similarity measure is a good candidate to provide good and robust recommendations in databases where there is a high degree of uncertainty about the validity of records. In fact, while the use of the BP similarity does not give substantially better recommendation prediction in databases like MovieLens and Fine Foods, its performance is consistently higher in the case of Jester Jokes.

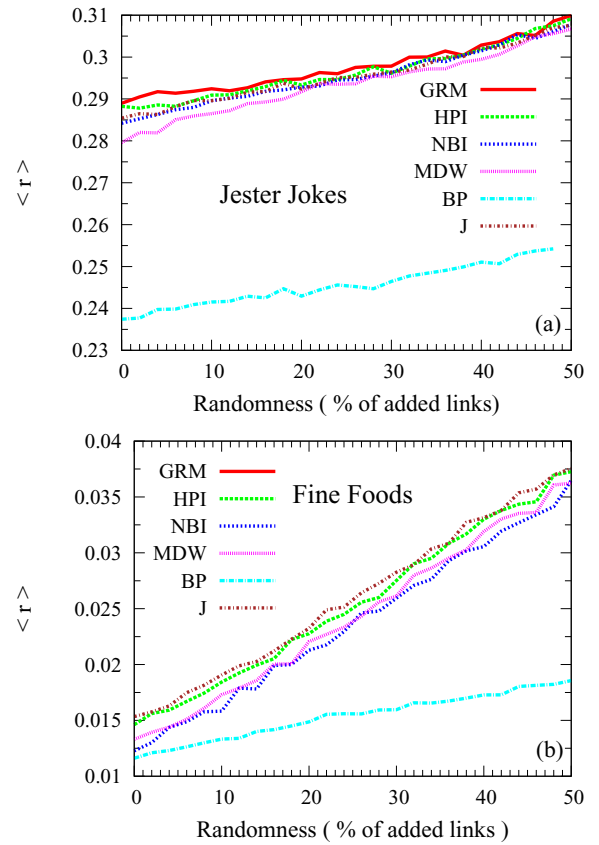


FIG. 6. (Color online) Mean validation values for some different methods as a function of the randomness with the Jester Jokes (a) and Fine Foods (b) databases.

VI. DIMENSIONALITY REDUCTION

The techniques of recommendation often make use of filters with the goal to reduce the role of databases' inaccuracies.




FIG. 7. (Color online) Average rank quality index $\langle r \rangle$ for the different recommendation methods as a function of the reduced matrix rank k . The mean is computed on $N_{Exp} = 10$ different realizations at $\lambda = 0$. The reduction is here done on the incidence matrix of the data set MovieLens. The Dimensionality Reduction method slightly improves the predictions at low k values.




FIG. 8. (Color online) Average rank quality index $\langle r \rangle$ for the different recommendation methods as a function of λ parameter with reduced matrix rank $k = 15$. The mean is computed on $N_e = 10$ different realizations using the Movie Lens data set.

In some cases, the filter is done by setting a threshold, as in the *correlation-thresholding* and *best-n-neighbors*, in order to keep the most relevant similarities and disregard the others [14,15]. In other cases, it might be convenient to manipulate the data matrices with more sophisticated techniques, such as the singular value decomposition (SVD), where the rank of the matrix (incidence or rating) is reduced by keeping the largest k eigenvalues and putting to zero the others. Following the approach of Sarwar *et al.* [16], we applied the SVD method to the MovieLens database and studied the performances of the different similarity measures. Figure 7 shows the average rank quality index $\langle r \rangle$ as a function of the rank k of the resulting reduced matrix. The index $\langle r \rangle$ is a slow monotonic increasing function of k . For low values of k , the performance is slightly improved for almost all the similarity measures with respect to the complete matrix, whose values are indicated with the arrows on the right edge of the plots. The measures confirm that dimensionality reduction can improve recommendation predictions.

We have also investigated, as in Sec. IV, the tradeoff between user similarities and object similarities. Specifically, we kept constant the rank of the matrix at a value of 15, and varied the parameter λ as before. Figure 8 shows the obtained results. It is apparent that the relative order relationship between the predictions of the different similarity measures is maintained. In other words, the method of dimensionality reduction of the adjacent matrix may improve the performances of the recommendations, for a set of suitably low k values with the different similarities, but its effect appears to be homogenous across the different similarity measures.

VII. CONCLUSIONS

We have considered three real-world users-items bipartite networks, we have investigated the performance of several traditional recommendation methods recently presented in the literature, and we have proposed two new similarity measures which take into account the heterogeneity of users and objects degrees. We showed that these two new similarity indexes can outperform traditional recommendation systems in most of the cases, even if there is a clear dependence of the results on the structural characteristics of the data set under study.

Then we focused on hybrid recommendation systems based on the convex combination of the recommendation scores induced by the similarity between users and objects, parametrized by a coefficient λ . We showed that different outcomes can be obtained in personalized recommendation methods by using similarity between users, or between objects, or a combination of the two. In some cases, the quality of recommendation as measured by the average rank quality index r is a convex function of the parameter λ . This means that the combination of different recommendation scores might actually provide better performance with respect than the employment of user or object similarities alone and, more importantly, that depending on the data set at hand, the quality of recommendation can be actually optimized through an appropriate tuning of λ . Conversely, for some similarity measures we observed a monotonically decreasing dependence of $\langle r \rangle$ on λ , so that the best recommendation is obtained by using an object-based similarity. We also investigated the robustness of recommendation systems to the addition and rewiring of edges, and the results suggested that the BP correlation similarity can consistently outperform other similarity measures in noisy data sets. Finally, we performed dimensionality reduction on the incidence matrix by making use of the singular value decomposition method. Also this last test confirmed the good performances of the proposed similarity measures.

In conclusion, although we do not observe a specific recommendation method outperforming all the others *in all conditions* and for all the data sets considered, it seems that recommendations based on MDW and BP are in general able to produce better results than those using other similarity measures. However, our results show that the performance of the recommendation methods depends on both the specific investigated database and on the way similarities between users and objects are used to derive recommendation scores.

ACKNOWLEDGMENTS

This work is partially supported by the EPSRC project GALE EP/K020633/1. This research utilized Queen Mary's MidPlus computational facilities, supported by QMUL Research-IT and funded by EPSRC grant EP/K000128/1.

- [1] G. Adomavicius and A. Tuzhilin, Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions, *IEEE* **17**, 734 (2005).
- [2] L. Lü, M. Medo, C.-H. Yeung, Y.-C. Zhang, and Z.-K. Zhang, Recommender systems, *Phys. Rep.* **519**, 1 (2012).
- [3] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, Item-based collaborative filtering Recommendation algorithms, in *Proceedings*

of 10th WWW International Conference, 1-5 may 2001 (ACM, NY, 2001), pp. 285–295.

- [4] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry, Using collaborative filtering to weave an information tapestry, *Commun. ACM* **35**, 61 (1992).
- [5] J. B. Schafer, D. Frankowski, J. Herlocker, and S. Sen, Collaborative filtering recommender systems, in *The Adaptive*

- Web*, edited by P. Brusilovsky, A. Kobsa, and W. Nejdl (Springer-Verlag, Berlin, Heidelberg, 2007), p. 291.
- [6] Y.-C. Zhang, M. Blattner, and Y.-K. Yu, Heat conduction process on community networks as a recommendation model, *Phys. Rev. Lett.* **99**, 154301 (2007).
- [7] T. Zhou, R. Q. Su, R. R. Liu, L. L. Jiang, B. H. Wang, and Y.-C. Zhang, Accurate and diverse recommendations via eliminating redundant correlations, *New J. Phys.* **11**, 123008 (2009).
- [8] T. Zhou, J. Ren, M. Medo, and Y.-C. Zhang, Bipartite network projection and personal recommendation, *Phys. Rev. E* **76**, 046115 (2007).
- [9] L. Lü and W. Liu, Information filtering via preferential diffusion, *Phys. Rev. E* **83**, 066119 (2011).
- [10] T. Zhou, Z. Kuscsik, J.-G. Liu, M. Medo, J. R. Wakeling, and Y.-C. Zhang, Solving the apparent diversity-accuracy dilemma of recommender systems, *Proc. Natl. Acad. Sci. USA* **107**, 4511 (2010).
- [11] T. Q. Qiu, Z.-K. Zhang, and G. Chen, Information filtering via a scaling-based function, *PLoS ONE* **8**, e63531 (2013).
- [12] X. Zhu, H. Tian, and S. Cai, Personalized recommendation with corrected similarity, *J. Stat. Mech.: Theor. Exp.* (2014) P07004.
- [13] K. Pearson, Note on regression and inheritance in the case of two parents, *Proc. R. Soc. Lond.* **58**, 240 (1895).
- [14] J. L. Herlocker, J. A. Konstan, A. Borchers, and J. Riedl, An algorithmic framework for performing collaborative filtering, in *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (ACM, NY, 1999), pp. 230–237.
- [15] Y. H. Cho and J. K. Kim, Application of Web usage mining and product taxonomy to collaborative recommendations in e-commerce, *Expert Sys. Appl.* **26**, 233 (2004).
- [16] B. M. Sarwar, G. Karypis, J. A. Konstan, and J. T. Riedl, Application of Dimensionality Reduction in Recommender System-A Case Study, in *Proceedings of the 2000 ACM WEBKDD Workshop* (ACM, NY, 2000).
- [17] L. Lü and T. Zhou, Link prediction in complex networks: A Survey, *Physica A* **390**, 1150 (2011).
- [18] P. Jaccard, The distribution of the flora in the alpine zone, *New Phytol.* **11**, 37 (1912).
- [19] R.-R. Liu, C.-X. Jia, T. Zhou, D. Sun, and B.-H. Wang, Personal recommendation via modified collaborative filtering, *Physica A* **388**, 462 (2009).
- [20] Q. Guo, R. Leng, K. Shi, and J. G. Liu, Heat conduction information filtering via local information of bipartite networks, *Eur. Phys. J. B* **85**, 286 (2012).
- [21] M. Tumminello, S. Miccichè, L. J. Dominguez, G. Lamura, M. G. Melchiorre, M. Barbagallo, and R. N. Mantegna, Statistically validated networks in bipartite complex systems, *PLoS ONE* **6**, e23377 (2011)
- [22] P. Resnick, N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl, GroupLens: An open architecture for collaborative filtering of netnews, in *Proceedings of the 1994 ACM Conference on Computer Supported Cooperative Work* (ACM Press/Addison-Wesley, New York, 1994), p. 175.
- [23] U. Shardanand and P. Maes, Social information filtering: algorithms for automating word of mouth, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (ACM Press/Addison-Wesley, New York, 1995), p. 210.
- [24] V. Hatzopoulos, G. Iori, R. N. Mantegna, S. Miccichè, and M. Tumminello, Quantifying preferential trading in the e-MID interbank market, *Quant. Fin.* **15**, 693 (2015).
- [25] See <http://www.grouplens.org/node/73>, <http://eigenstate.berkeley.edu/dataset/>, and <https://snap.stanford.edu/data/web-FineFoods.html>, respectively.