



Testing the Impact of Semantics and Structure on Recommendation Accuracy and Diversity

Pedro Ramaciotti Morales, Lionel Tabourier, Raphaël Fournier-S’Niehotta

► To cite this version:

Pedro Ramaciotti Morales, Lionel Tabourier, Raphaël Fournier-S’Niehotta. Testing the Impact of Semantics and Structure on Recommendation Accuracy and Diversity. ASONAM 2020. IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, 2020, La Hague (virtual), Netherlands. <hal-02999863>

HAL Id: hal-02999863

<https://hal.science/hal-02999863v1>

Submitted on 11 Nov 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L’archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d’enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Testing the Impact of Semantics and Structure on Recommendation Accuracy and Diversity

Pedro Ramaciotti Morales¹, Lionel Tabourier², and Raphaël Fournier-S'niehotta³

¹Médialab, Sciences Po, Paris, France

²Sorbonne Université, CNRS, LIP6, F-75005 Paris, France

³CEDRIC, CNAM, Paris, France

Abstract

The Heterogeneous Information Network (HIN) formalism is very flexible and enables complex recommendations models. We evaluate the effect of different parts of a HIN on the accuracy and the diversity of recommendations, then investigate if these effects are only due to the semantic content encoded in the network. We use recently-proposed diversity measures which are based on the network structure and better suited to the HIN formalism. Finally, we randomly shuffle the edges of some parts of the HIN, to empty the network from its semantic content, while leaving its structure relatively unaffected. We show that the semantic content encoded in the network data has a limited importance for the performance of a recommender system and that structure is crucial.

1 Introduction

A Recommender System (RS) helps users in selecting items from large sets when their size makes individual consideration impractical or even impossible [17]. To propose relevant recommendations to users, a RS may process previous choices made by users, item/user meta-data, and contextual information. Initially, the performances of RS were measured with the so-called utility or accuracy metrics [6],

which evaluate, for example, the error committed when predicting the rating a user would give to a specific item. Later, other properties of recommendations beyond accuracy were suggested, because users also wish for surprising recommendations and thus should be proposed diverse items [12]. Moreover, diversifying recommendations is also desirable to guarantee that people are not shut in a small subset of homogeneous recommendations, a phenomenon known as a filter bubble [15].

RS models are often separated between content-based filtering and collaborative filtering approaches. Many recent RS tend to use hybrid methods that mix both approaches [1]. Among these, methods relying on the Heterogeneous Information Network (HIN) framework are gaining popularity [24]. HINs are graph-based representations with different types of node: some represent users, others represent items, and a new type of node (or layer) is added for each kind of information in the system [2]. Then nodes of different types are connected with links representing a specific type of relation, such as *a user selected an item*, or *an item belongs to a given genre*, etc. A simple example consists in *readers* linked to *books* they read, which are in turn linked to *genres* (categories).

In this work, we want to evaluate how different parts of a HIN impact the diversity and the accuracy of the recommendations a RS based on this HIN may produce. Using a RS derived from Pathsim [24], which relies on random walks on HINs, we may gradually include different parts of the HIN into the RS in a controlled way. These different parts encode different meta-data on users and items: some nodes may represent the demographic group of users (say, their gender), others may represent a category of items (e.g., drama or comedy for movies), etc. Most current diversity measures may not benefit from the full structural information of networks, since they are defined only with the output of the recommendation process (the lists of recommended items and their features) [9, 22]. Thus, our evaluation of diversity relies on a framework especially suited to HINs [16].

Another important goal of our work is to measure whether the performance of a HIN-based RS is mostly due to the nature of the semantic information encoded, or to a *structural shortcut effect*. A HIN representation may contain connections between items or users who may not share a lot of common features. For instance, two users might have different tastes while living in the same city and a HIN representation would connect them through their zip code. A HIN-based RS would use these “shortcuts” to recommend items otherwise unreachable, affecting diversity. We address this problem of balance between structural and semantic information by using a *configuration model* [14] of our HIN. Randomly shuffling edges in a given part of the HIN (while keeping the degree distribution unaffected), we separate the effects of those semantic and structural information.

To summarize, the research questions addressed are:

- Q1: How to test the effect of different parts of a HIN on the accuracy and the diversity of recommendations?, and
- Q2: Are the effects of different parts of a HIN on accuracy and diversity due to the semantic content encoded in them?

This article is structured as follows. In Section 2, we discuss the related work. In Section 3, we provide the definitions of concepts used in this paper, in particular the formal definition of a HIN and suited diversity measures. We also expose the design of our HIN-based RS model. In Section 4, we detail a series of experiments which correspond to one of our contributions:

1. we provide the first application of *network diversity measures* to recommendations computed on HINs (Sec. 4.1).
2. we use a simple, well-known procedure to compute recommendations on HINs [24], and test it on the datasets to answer question Q1. We show that depending on the information included, such a method can yield accurate and diverse recommendations (Sec. 4.2).
3. finally, we use edge shuffling of some parts of a HIN to show that it is not always possible to explain the performance of recommendations by the semantic information encoded in HINs, thus answering question Q2 (Sec. 4.3). To the best of our knowledge, this is the first work to investigate how the global structure of a HIN affects accuracy and diversity of recommendations.

2 Related Work

HIN-based RSs are rapidly developing into a subdomain of information filtering and recommendation [2, 19]. HINs enable mixing past choices of users with semantic information about items and users in a unified graph-based representation. In a HIN, if a set of users are linked to books they have chosen, and these books linked to genres, for example, a HIN-based RS may exploit the collective choices made by users (the set of links representing choice), and the genres of books chosen by users. Recommendations based on content (recommending books of previously chosen genre), but also on collective choice (recommending books chosen by users that have similar taste) are then possible. The HIN formalism has

been shown to perform well and to improve interpretability [18, 21]. Moreover, HINs can deal with cross-domain recommendation, often useful in cold-start situations [25], and can exploit social relations between users [27]. They also allow predictions based on implicit feedback [28].

RS evaluation usually focuses on *accuracy*: measuring the error on predicted ratings with metrics such as RMSE or MAE, or measuring the error committed on items selected by users using metrics such as precision, recall, or F1-score. However, almost two decades ago, it was acknowledged that properties beyond accuracy play an important role in user satisfaction [6, 12]. These properties aim at grasping how novel or surprising a recommendation is [6] or how diverse the proposed recommendation list is [33]. There is little consensus on the naming of these desirable properties, but they may be collectively referred to as *diversity* measures, indicating a degree of (dis-)similarity between recommended items. Some authors consider two items similar if they are chosen by the same users [7, 31], while others consider items similar if they are of the same type [33]. These notions are often called *novelty* or *intra-list similarity*. Some works focus on the improbability, or the *surprisal*, of a recommendation considering the users' past choices, while others focus on *personalization*, *i.e.*, the degree to which recommendations are different from one user to another [32]. Measuring diversity is also important for specific applications, such as detecting context changes [10] or measuring phenomena related to diversity loss, such as filter bubbles and echo chambers [15]. An extensive survey recently compiled diversity measures for RS [9].

In information filtering, searches in graphs sometimes call to some notion of diversity on node rankings (*i.e.*, item rankings) [11, 26]. While many studies consider diversity with respect to some node labeling, few exploit the variety of paths formed by relations between nodes of different types [9]. Dubey et al. [3] use, as we do in this article, the notion of random walk along these paths, which are usually called *meta-paths* in HINs. Furthermore, some works use the notion of random walk in relation to diversity [8, 13], although not as a measure to evaluate resulting recommendations of RSs in HINs. While there are theoretical propositions of diversity measures adapted to HIN [16], there are no available experiments that illustrate their applicability.

3 Heterogeneous Information Networks and Diversity

In this section we give definitions and notations about HINs that are used in the rest of this paper. By combining these definitions with random walks on graphs, we are able to define the diversity measures that are then exploited in Section 4.

3.1 Heterogeneous Information Networks

Definition 1 (Heterogeneous Information Network [20]). *A Heterogeneous Information Network (HIN) $\mathcal{H} = (\mathcal{G}, \mathcal{A}, \mathcal{R}, \varphi, \psi)$ is a directed graph $\mathcal{G} = (V, E)$ with mapping functions $\varphi : V \rightarrow \mathcal{A}$ and $\psi : E \rightarrow \mathcal{R}$, where \mathcal{A} and \mathcal{R} are respectively node and edge label sets.*

Slightly simplifying notations, we refer to a group of nodes or edges by its label. Edge labels in \mathcal{R} represent relations between entities of different object groups in \mathcal{A} .

Definition 2 (Heterogeneous Information Network Schema [20]). *The network schema of a HIN $\mathcal{H} = (\mathcal{G}, \mathcal{A}, \mathcal{R}, \varphi, \psi)$ is a directed graph, the nodes of which are the types \mathcal{A} , and the edges are the relations of \mathcal{R} .*

A HIN schema (or simply *schema*) is a graphical and summarized representation of a HIN, which shows the groups of nodes and links. Figure 1 shows a basic example of HIN schema, corresponding to a platform where users would rate items and items belong to predefined categories. Other examples can be found in Section 4.1.




Figure 1: HIN schema of a simple case where users (U) rate items (I) and items belong to types (Ty), thus defining edge labels R_{rates} and R_{Ty} .

Definition 3 (Link Group). *A link group, represented by an edge in the HIN schema $R \in \mathcal{R}$, is the subset of edges in E that are of type R , linking source object group $S \in \mathcal{A}$ and target object group $T \in \mathcal{A}$. It is denoted simply by R , and its inverse (made of the reversed edges of a link) by R^{-1} .*

Definition 4 (Meta-path [20]). A meta-path of a HIN is a path on its schema, and it is denoted by $\Pi = A_1 \xrightarrow{R_1} A_2 \xrightarrow{R_2} \dots \xrightarrow{R_l} A_{l+1}$, or more simply by $\Pi = R_1 R_2 \dots R_l$.

In the example of Figure 1, $\Pi = R_{rates} R_{Ty}$ is a meta-path connecting users to the types of items that they have rated.

3.2 Random Walks and Meta-Path Diversities on HINs

Given a HIN $\mathcal{H} = (\mathcal{G}, \mathcal{A}, \mathcal{R}, \varphi, \psi)$ let us consider a random walk along a meta-path Π starting in $S \in \mathcal{A}$ and ending in $T \in \mathcal{A}$. Let us also consider the random variable X , the node of T where the random walk along Π ends. We are interested in the probability mass function (PMF) over the nodes of T , denoting p_Π the PMF $P(X = t \in T)$ when the walk started randomly at any node of S . Similarly, let us denote by $p_\Pi(s)$ the PMF $P(X = t \in T | s \in S)$, when the walk started at the node $s \in S$. To illustrate this better, let us consider a HIN with 2 users rating 3 items classified in 2 types, corresponding to the schema in Figure 1. In Figure 2, we illustrate the computation of PMFs for such random walks on the simple example presented in Figure 1, following the meta-path $\Pi = R_{rates} R_{Ty}$: p_Π when starting randomly in a set of users $U = \{u_1, u_2\}$, and $p_\Pi(u_2)$ when starting from user u_2 .




Figure 2: Example of a HIN (left) having a schema as in Figure 1, with 2 users rating 3 items classified into 2 types, and computation of p_Π (middle) and $p_\Pi(u_2)$ (right) for meta-path $\Pi = R_{rates} R_{Ty}$.

Different diversity measures can then be computed from these PMFs (e.g., Shannon entropy, Gini coefficient). We use *perplexity*, denoted by \mathcal{P} and computed for a PMF p as $\mathcal{P}(p) = 2^{E(p)}$, with $E(p) = -\sum_{i=1}^{|T|} p_i \log_2(p_i)$, where $E(p)$

is the Shannon entropy of the PMF. We select perplexity because of its intuitive interpretation: it measures the degree of unexpectedness of a random variable. Perplexity accounts for *variety* (number of possibilities available) and *balance* (how probability is distributed among those possibilities) [23]. Interestingly, the perplexity of a PMF $p = (p_1, p_2, \dots)$ can be interpreted as the number of elements of a uniform PMF that has the same Shannon entropy as p . For instance, $p = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ has perplexity $\mathcal{P}(p) = 3$ (i.e., uniform over 3 elements) while $p' = (\frac{6}{10}, \frac{2}{10}, \frac{1}{10}, \frac{1}{10})$ has perplexity $\mathcal{P}(p') = 2.97$, i.e., p' has the same entropy as a uniform distribution over 2.97 elements. So p' is a distribution over four possibilities, but has lower diversity than p .

Using these notions, and following [16], we define the two following diversity measures for a meta-path Π of a HIN \mathcal{H} :

$$\begin{aligned} \text{Mean Individual diversity: } \mathcal{P}_{\text{MI}}(\Pi) &= \left(\prod_{s \in S} \mathcal{P}(p_{\Pi}(s)) \right)^{1/|S|}, \\ \text{Collective diversity: } \mathcal{P}_{\text{Col}}(\Pi) &= \mathcal{P}(p_{\Pi}), \end{aligned}$$

with \mathcal{P} being the perplexity.

The Mean Individual (MI) diversity along a meta-path Π is the geometric mean of the perplexities of the PMFs corresponding to random walks starting on each element $s \in S$ and ending in T . It can thus be interpreted as the diversity in terms of elements of T available on average for one element of S . Collective (Col) Diversity along meta-path Π is the perplexity of the PMF corresponding to a single random walk starting randomly on an element $s \in S$ and ending in T . It may be interpreted as the diversity in terms of elements of T available to the entire group of elements of S . In the example of Figure 2, the collective diversity of types available to users is $\mathcal{P}_{\text{Col}}(\Pi) = \mathcal{P}(3/4, 1/4) = 1.7547$, and the mean individual diversity is

$$\mathcal{P}_{\text{MI}}(\Pi) = (\underbrace{\mathcal{P}(1, 0)}_{p_{\Pi}(u_1)} \cdot \underbrace{\mathcal{P}(1/2, 1/2)}_{p_{\Pi}(u_2)})^{\frac{1}{2}} = (1 \cdot 2)^{\frac{1}{2}} = \sqrt{2}.$$

3.3 Recommendations as Processes on HINs

A recommendation process consists in proposing items to a user that he or she has not selected before. In a HIN framework, it amounts to creating new edges:

Definition 5 (Recommendation on a HIN). *A recommendation process on a HIN $\mathcal{H} = (\mathcal{G}, \mathcal{A}, \mathcal{R}, \varphi, \psi)$ is a process $F(\mathcal{H})$ that produces a new collection of edges $F : \mathcal{H} \mapsto E_{rec} \in V \times V$. The edge set E_{rec} is composed of links between users and their recommended items.*

A HIN can then be updated with a new relation R_{rec} such that $\psi(e) = R_{rec}$ for $e \in E_{rec}$, so that recommended edges can be used in meta-paths. In the following section, we are interested in including recommendation edges in meta-paths to measure the diversity related to these paths, as described in Subsection 3.2. We may then measure how *diverse* the recommendation process is.

3.4 A HIN Graph Spreading Recommender System

In order to produce recommendations that can include different parts of a HIN in a controlled and interpretable fashion, we consider a recommendation procedure well-known to the RS community (cf. [24]). This method belongs to a larger family of RS based on random walks (e.g., [8, 13]).

Precisely, it consists in computing the probability produced by random walks along a meta-path, starting on a single user and ending on the set of items. Then, we simply recommend the items with the highest probability that have not been previously chosen. This allows us to measure the effect of different parts of a HIN in an interpretable way. Indeed, we can consider different meta-paths that include different meta-data (e.g., book genre, user location), and produce different probabilities that can be added with different coefficients to modulate their relative importance in recommendations. For instance, we may give a high weight to the meta-path related to demographic information about the users and a light weight to information related to items typologies, or vice versa. Let us formalize this type of recommendation by considering K different meta-paths Π_k (with $k = 1, \dots, K$) starting on $S \in \mathcal{A}$ and ending in $T \in \mathcal{A}$. Let us consider the parameters $\alpha_1, \dots, \alpha_K$, such that $\sum \alpha_k = 1$ and $0 \leq \alpha_k \leq 1$. For every element $s \in S$ (typically the user group), we compute a score \mathcal{S} as a PMF over T (typically the item group):

$$\mathcal{S}(t|s) = \sum_{k=1}^K \alpha_k p_{\Pi_k}(s), \text{ for } t \in T. \quad (1)$$

Using the score $\mathcal{S}(t|s)$ we recommend the top ranking elements of T for $s \in S$ (among those not already selected). This RS shares some principles with the one detailed by Zhou et al. [32], in which mass (called *resource* in [32]) is assigned

to items and then distributed through edges to find relevant items according to the final resource distribution.

It should be noticed that this RS is not new and our point is not to argue that it performs better than other methods. However, we show in the next section that its accuracy is comparable to that of popular methods. Its main interest here is that it allows us to include different parts of a HIN in the recommendation process to measure separately the effects on accuracy and diversity of these parts.

4 Experiments


Here we present the data used, the protocols for our experiments¹, and their results to answer our research questions.

4.1 Datasets and Meta-Path Diversities of HIN

We use two datasets in our experiments, the well-documented MovieLens 100K dataset (ML100K) [4] and the Douban Movie (DM) dataset [30]. The latter has been used in HIN-based RS studies, as it presents a rich HIN structure [21]. We detail in Figure 3 the characteristics of these datasets: groups of entities, their sizes, and the relations that join them. We also show the corresponding HIN schemas with their object groups and link groups (relations). Both datasets contain information on the rating (from 1 to 5) given to films by users, and the fact that a user has rated a movie is represented by the relation R_{rating} . We consider that a user likes a film if he or she rated it 3 or more (for both datasets), and we refer to the corresponding relation as R_{likes} . Figure 3 also shows the liking and recommendation relations with dashed lines, R_{likes} and R_{rec} , the latter indicating that a RS may propose a recommendation of an item to a user.

¹The codes and data used for these experiments are available at <https://github.com/pedroramaciotti/HINPy>

MovieLens 100K (ML100K)					
Relation	Start Obj. Group	Size	End Obj. Group	Size	Links
R_{Ty}	movie	1664	type	19	2863
R_{Ye}	movie	1664	release	72	1664
R_{rates}	user	943	movie	1664	99965
R_{Oc}	user	943	occupation	21	943
R_{Ag}	user	943	age group	61	943
R_{Ge}	user	943	gender	2	943
R_{Lo}	user	943	location	795	943
R_{likes}	user	943	movie	1664	82495



Douban Movie (DM)					
Relation	Start Obj. Group	Size	End Obj. Group	Size	Links
R_{Ac}	movie	6971	actor	3004	15584
R_{Di}	movie	6971	director	789	3314
R_{Ty}	movie	6971	type	36	15598
R_{Gr}	user	3022	usergroup	2269	7054
R_{Lo}	user	3022	location	244	2491
R_{rates}	user	3022	movie	6971	195493
R_{Fr}	user	3022	user	3022	779
R_{likes}	user	3022	movie	6971	182069




Figure 3: Summary and HIN schema of the MovieLens 100K (top) and Douban Movie (bottom) datasets.

Now, we can illustrate the use of the diversity measures on HINs defined in Section 3 with the ML100K and DM datasets. While diversity measures from the literature essentially rely on how proposed or selected items are distributed into types, the diversity measures described in Section 3.2 allow to consider a larger spectrum of diversity notions. Let us consider the example of a diversity computed on meta-path $\Pi_1 = R_{\text{rec}}R_{\text{Ty}}$. It is the diversity of types of movies in a recommendation list, consequently it captures the same intuition as the *Intra-List Similarity* does in [33]. However, the meta-path based diversity allows us to express other forms of diversities. For instance, the meta-path $\Pi_2 = R_{\text{rec}}R_{\text{likes}}^{-1}R_{\text{likes}}R_{\text{Ty}}$ relates to the diversity of types of movies liked by users who also liked movies that they were recommended.

In Figure 4, we report experimental diversity measures on both datasets. We consider meta-paths of length 3 starting in S and ending in T (with $S, T \in \mathcal{A}$), which are node groups containing semantic information about users and items respectively, so $\Pi = R_S^{-1}R_XR_T$. Here X stands for *rating*, *liked* or *recommended*. Then, we compute the Mean Individual diversity on these paths: $\mathcal{P}_{\text{MI}}(\Pi)$. Regarding the recommendation links, we use two standard RSs in these experiments, each recommending a list of 5 items to the users: *User-Based Collaborative Filtering* (UBCF) and *Implicit Pure Popularity* (IPP), IPP recommends to the users the top-5 most liked movies, not already rated by the user. The figure reads as follows: it reports Mean Individual diversities for all combinations of node groups S and T and R_X is either R_{likes} , $R_{\text{rec}}(\text{UBCF})$ or $R_{\text{rec}}(\text{IPP})$. For instance, $\mathcal{P}_{\text{MI}}(R_{\text{Lo}}^{-1}R_{\text{rec}}(\text{UBCF})R_{\text{Ty}})$ is the Mean Individual diversity of movie types recommended to users using UBCF, averaged over locations.

The measurements of these diversities provide insights about the structure of the HIN and about the RS. First, notice that most $\mathcal{P}_{\text{MI}}(\Pi)$ using *recommendations* are lower than those using *likes*. This is only a consequence of the fact that recommended lists are of length 5, while the average user likes larger numbers of items in these datasets. Similarly, the fact that $\mathcal{P}_{\text{MI}}(R_S^{-1}R_XR_{\text{Ac}})$ is much larger than $\mathcal{P}_{\text{MI}}(R_S^{-1}R_XR_{\text{Ty}})$ or $\mathcal{P}_{\text{MI}}(R_S^{-1}R_XR_{\text{Di}})$ on DM dataset reflects the fact that movies have more actors than directors or types/genres. Other phenomena are revealed by these diversity measures. For example, the diversity of release year of films by age groups (*i.e.*, of $R_{\text{Ag}}^{-1}R_{\text{likes}}R_{\text{Ye}}$) in the ML100K dataset is low. This could point to the fact that users from a given age group tend to choose films from a narrow period of time, however testing this hypothesis falls outside the scope of this article. Note also that these diversity measures allow us to compare diversities from one RS to another. For example UBCF produces recommendations that are consistently less diverse than those produced using IPP on the ML100K dataset.




Figure 4: Mean individual diversities of different meta-paths $\Pi = R_S^{-1} R_X R_T$ for ML100K (a) and DM (b) datasets. S are indicated in rows, T in columns, and X at the bottom for each group of diversities (5 items recommended per user).

4.2 Effect of the HIN Structure on Accuracy and Diversity of Recommendations

Having established and illustrated the use of meta-path diversities, we now focus on the effect that the structure of HINs has on accuracy and diversity of recommendations. For this purpose, we use the RS described in Section 3.4. It allows us to include in the recommendation process different parts of a HIN in a controlled way. Moreover, its underlying mechanic, based on meta-path constrained random walks, is common to several techniques of HIN-based recommendation (e.g., [21, 24, 29]). Note that our goal here is to explore the consequences on accuracy and diversity of integrating different parts of a HIN, rather than optimizing this RS to obtain the best possible performance.

Many choices are available, as we can mix all kinds of meta-paths to produce a recommendation. However, we choose to consider a RS that mixes only two meta-paths for each dataset: one denoted $\Pi(X)$ which includes a *user-content* object group (e.g., the location of a user), and one $\Pi(Y)$ including an *item-content* object group (e.g., the genre of a movie). By tuning the balance between these two meta-paths, we can control what kind of information is taken into account to produce a recommendation. Practically, Equation (1) here takes the form:

$$\mathcal{S}(t|s) = (1 - \alpha) (p_{\Pi(X)}(s)) (t) + \alpha (p_{\Pi(Y)}(s)) (t)$$

$s \in \mathcal{U}$ and $t \in \mathcal{I}$ and parameter $\alpha \in [0, 1]$ tunes the balance between $\Pi(X)$ and $\Pi(Y)$. So, $\alpha = 1$ means that only the item-content is used in computing recom-

mendations and when $\alpha = 0$ only the user-content is used. Precisely, the user-content path $\Pi(X)$ is $R_X R_X^{-1} R_{\text{likes}}$ where X stands for a user-content group, and the item-content path $\Pi(Y)$ is $R_{\text{likes}} R_Y R_Y^{-1}$ where Y stands for an item-content group. Figure 5 illustrates the meta-paths selected on the HIN schema.




Figure 5: Schematic illustration of the recommendation procedure used in the experiments, where parameter α modulates two meta-paths, $\Pi(X)$ and $\Pi(Y)$, containing user and item meta-data.

We use $\alpha \in \{1.0, 0.8, 0.6, 0.4, 0.2, 0.0\}$ to show the effect of gradually moving from recommendations purely based on item meta-data ($\alpha = 1$) to recommendations based purely on user meta-data ($\alpha = 0$). While there are several possible combinations of meta-paths (8 combinations for ML100K and 6 for DM), we only focus on a few of them (4 for each dataset) for ease of illustration and also because they suffice to show the effects that the HIN structure has on accuracy and diversity.

4.2.1 Accuracy

Figures 6a and 6b show the F1-score of recommendations using the ML100K and DB datasets respectively. For each experiment, we set a list size (the number of items included in a recommendation list proposed to each user). For a given list size and α , a random 10% of user-item links are hidden. We compute the items recommended to users (corresponding to the edges R_{rec} in Figure 3) with our HIN-based RS and compare it to the test set to evaluate the F1-score of our prediction. Each figure in the mosaic corresponds to a combination of meta-paths, with a given X and a given Y .

While it is not our main objective to characterize the performance of this recommendation method, it is worth noticing that it may produce relevant recommendations: some combinations of meta-paths and α result in F1-scores comparable to other RS in different settings (see, *e.g.*, [5, 18]). Note also that whatever the

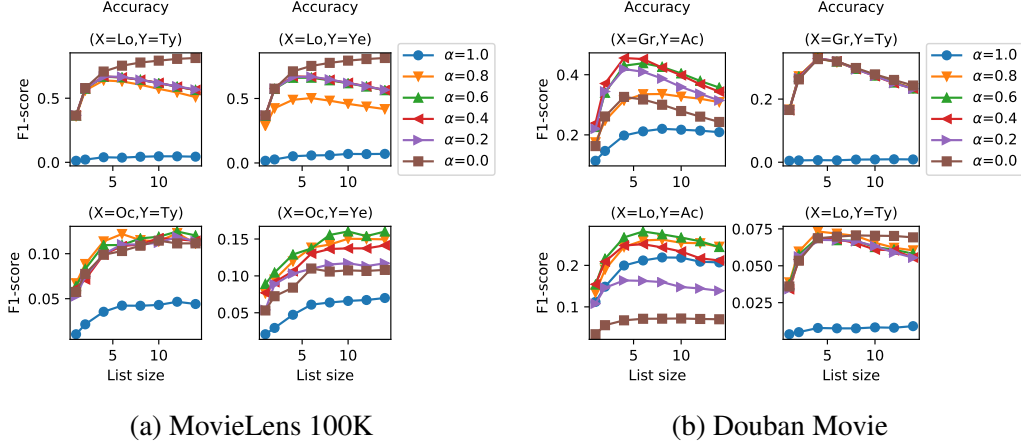


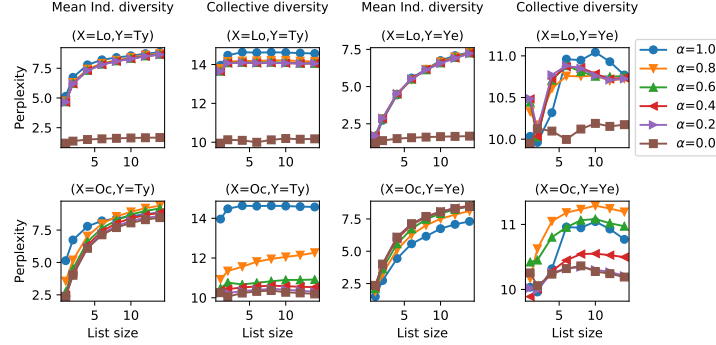
Figure 6: F1 score for different list sizes of recommended items and different values of α . The meta-paths used are $\Pi(X)$ and $\Pi(Y)$, where X is either Gender or Location (in rows) and Y is either Actors or Type (in columns).

list size, the pure item-content recommendation is never the highest F1-score. In other words, accuracy is, unsurprisingly, increased (up to a certain point) by the addition of user-content meta-data for both datasets.

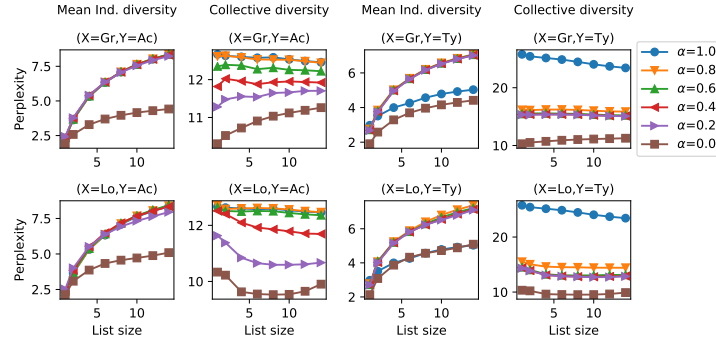
On Douban Movie with a fixed list size, the best F1-scores are obtained with a mixture of item- and user-related information. This is less true for ML100K, where user-related information seems often more valuable than item-related information. We also notice that the range of F1-score values varies significantly depending on the balance between meta-paths: we see variations by up to a factor of 14 on ML100K (list of size 15 using user locations and movie types). This observation is consistent with the idea that while some meta-data brings relatively irrelevant information for recommendation (e.g. movie year of release in ML100K), other are much more valuable for the accuracy of the recommendation (e.g. user location in ML100K).

4.2.2 Diversity

For the same recommendation protocol, we now measure the Mean Individual and Collective diversities of the recommendation lists, in terms of types of items recommended. In other words, we compute $\mathcal{P}_{MI}(\Pi)$ and $\mathcal{P}_{Col}(\Pi)$ with the $\Pi = R_{rec}R_{Ty}$ meta-path. The results are displayed in Figures 7a and 7b, following the same display organization as figures in 4.2.1.



(a) MovieLens 100K



(b) Douban Movie

Figure 7: Collective and Mean Individual diversities in terms of types of items recommended to users, for different list sizes and recommendation meta-paths modulated by α , for MovieLens 100K dataset (a), and Douban Movie dataset (b).

Figures 7a and 7b highlight the interpretation associated with diversity measures based on meta-paths, and especially the contrast between Mean Individual diversity and Collective diversity. For both datasets, Mean Individual diversity monotonically increases with the list size, which is expected: more items recommended to a user implies that the user may access more types. In other words, *variety* increases with the list size (although not necessarily *balance*). However, this tendency is limited as perplexity cannot be larger than the available number of types (19 in ML100K, 36 in DM). While the Mean Individual diversity grows with the number of recommended items, Collective diversity may stagnate or even decrease. For instance, we can read on DM dataset that Collective diversity de-

creases with the list size for all $\alpha > 0$ when $X = \text{Gr}$ (Groups of users) and $Y = \text{Ty}$ (movie Types). This result is subtle but important in our opinion. It means that even if each user is individually recommended more diverse items, the whole group of users are recommended less diverse items.

Other trends can be isolated from these experiments, we report a few of them. First, we observe that including user meta-data tends to improve accuracy, but also diminishes Collective diversity. This is more obvious on ML100K dataset, where combinations which bring higher F1 scores, are often the ones which have the lowest diversities (both Collective and Mean Individual). Also, including more user meta-data usually has little effect on Mean Individual diversity. This observation is quite surprising as user meta-data can connect users of the HIN who potentially have very little in common (for example users of a same gender), so we would expect that these links give access to a large variety of items. We think that these observations are certainly related to the HIN structure and we begin to address this question in the next section.

4.3 Testing Semantic Relevance of Links Using the Configuration Model

The previous results raise some interesting questions. Why does the inclusion of certain meta-paths changes the accuracy and the diversity more than others? What is the property of user-content links that allows for more accurate recommendations? A natural assumption is that they contain relevant semantic information about the origin of the users' preferences.

To test this hypothesis, we propose an experiment inspired by the field of social network analysis (SNA). To identify the specificities of a real network in SNA, it is usual to compare its structure to a random graph model with a similar degree sequence, that is to say a graph in which nodes have the same number of connections as in the original graph, but distributed randomly. This model is often referred to as the *configuration model* [14]. A usual way to perform such a comparison consists in iteratively shuffling the links of the existing graph by exchanging their extremities. Here, we propose to follow this process for the edges of a specific relation (link group) of the HIN. To illustrate this idea, let's consider the example of the user-content links R_{Lo} : by shuffling only those links, we do not modify the number of individuals per location. Thus we leave an essential property of the original HIN untouched, that is its degree sequence. However, we empty this part of the HIN of its semantic meaning, because the location associated to a user is

now random.

As for the protocol itself, we proceed as follows: we focus on a recommendation based on the mixing of two meta-paths, one using user-based information and another using item-based information, as we did previously. We select the combinations that produced the most accurate recommendations: user location (Lo) and movie type (Ty) for ML100K, and user groups (Gr) and movie actors (Ac) for DM. In addition to the recommendation produced on the original ML100K HIN (respectively DM HIN), which was reported in Section 4.2, we make recommendations on several realizations of a HIN where R_{Lo} (resp. R_{Gr}) have been randomly shuffled. Then we compare the accuracy, Mean Individual and Collective diversities obtained on the real HIN to the values on the randomized HINs. We better detail this procedure thereafter.

We produce 100 randomized versions of the HIN for the ML100K by shuffling the edges in the R_{Lo} link group, and 20 randomized HIN for the DM dataset by shuffling the edges in the R_{Gr} link group. Then, for each list size and for $\alpha \in \{0.8, 0.6, 0.4, 0.2, 0.0\}$, we compute the 0.1–0.9 quantile range of the F1-score, Mean Individual, and Collective diversities. Note that $\alpha = 1$ produces recommendations using only item-related meta-paths and is of no interest as it does not include randomized parts. Lower α values imply that randomized links from R_{Lo} (ML100K) or R_{Gr} (DM) are used in the recommendation process. A value of $\alpha = 0$ means that only the suffled link group R_{Lo} (ML100K) or R_{Gr} (DM) was used in computing recommendations. The resulting quantile ranges, compared to the values obtained with the original HIN from Section 4.2 are shown in Figure 8.

We observe in all cases that the recommendations on the original HIN and on its randomized versions are comparable in terms of accuracy and diversities. These results are counter-intuitive: it seems that emptying important parts of the HIN of its semantic content does not have a significant impact on the recommendation. These results are surprising as our recommendation process relies exclusively on the number of paths in the HIN connecting a user to an item. Our results are of course limited to the selected parameters and datasets used and call for further investigation in other contexts. Still, they tend to disprove the underlying hypothesis that the relevance of meta-path recommendations only comes from the semantic information that they contain. They could indicate that the number of connections of a node (which is not affected by the randomization) is the most – and even the only – important information in the recommendation process based on the HIN structure. This new hypothesis is not tested in this article as it demands a wider experimental setting than the one considered for the purposes stated in the

Dataset	User-content meta-path	Related user-content object group (to be shuffled)	Item-content meta-path	Related item-content object group	Number of randomizations
ML100K	$R_{Lo}R_{Lo}^{-1}R_{likes}$	Location (Lo)	$R_{likes}R_{Ty}R_{Ty}^{-1}$	Type (Ty)	100
DM	$R_{Gr}R_{Gr}^{-1}R_{likes}$	Usergroup (Gr)	$R_{likes}R_{Ac}R_{Ac}^{-1}$	Actor (Ac)	20

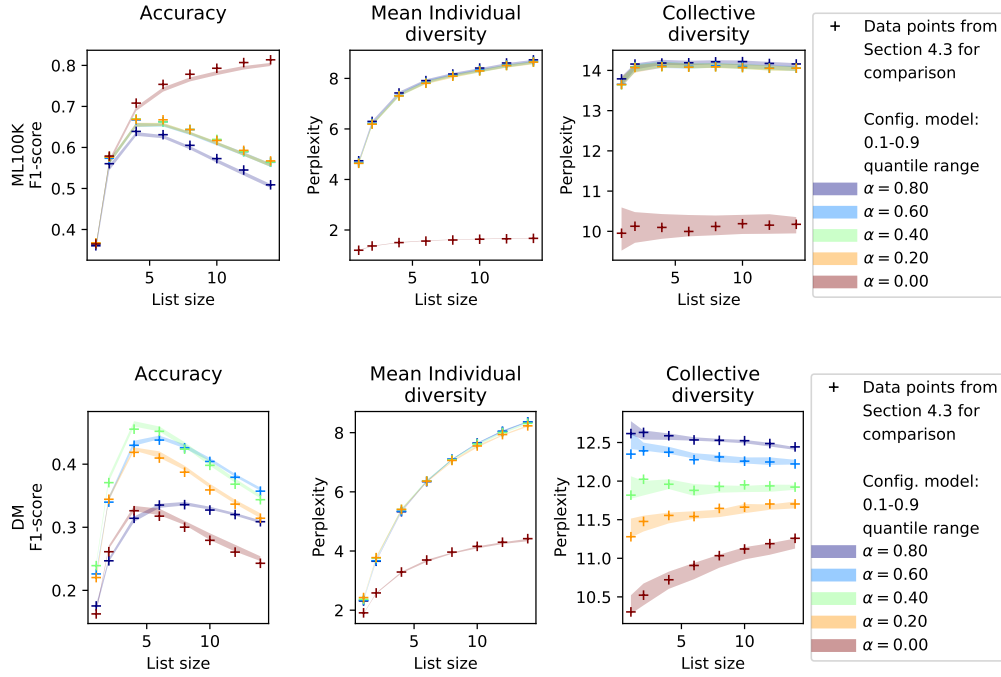


Figure 8: Comparison between recommendations on randomized HINs compared to the original HIN: 0.1-0.9 quantile ranges for F1-score, Mean Individual, and Collective diversities for different list sizes, α values.

introduction.

5 Conclusion

We used recently proposed diversity measures for the evaluation of recommendations on HINs. They are based on the *perplexity* of PMFs resulting from random walks along meta-paths over a Heterogeneous Information Network. While standard diversities measures usually focus on the distribution of the recommended

items into types, we may measure diversity on an individual and collective level, revealing that these two approaches may lead to different conclusions. We illustrated the experimental uses of these measures with an existing RS and two well-known datasets. We then presented a method to explore the effects of including different meta-paths in the recommendation process, evaluated its accuracy, and analyzed recommendations produced by combining different meta-paths. We showed how including different parts of a HIN in a recommendation may significantly change accuracy and diversity. For both datasets, including meta-paths with user-content information improved accuracy but often reduced collective diversity.

Finally, we tested the influence of the semantic content on the accuracy and diversity of recommendations. We randomized part of the HIN structure to remove its semantic meaning, while keeping the degree sequence of the HIN untouched. Surprisingly, this random shuffling did not change significantly accuracy or diversity. While these results are limited to a restricted number of datasets and parameters, they indicate that the effects of the information encoded in the structure of a HIN on recommendations are not only due to their semantic meaning.

Acknowledgements

We warmly thank Elizabeth Doggett for her proofreading of the manuscript. This work has been funded by the French National Agency of Research (ANR) under grants ANR-15-CE38-0001 (Algodiv), ANR-19-CE38-0006 (GOPI), and ANR LabCom FiT.

References

- [1] Robin Burke. Hybrid recommender systems: Survey and experiments. *User modeling and user-adapted interaction*, 12(4):331–370, 2002.
- [2] Robin Burke, Fatemeh Vahedian, and Bamshad Mobasher. Hybrid recommendation in heterogeneous networks. In *Proceedings of UMAP’14*, pages 49–60. Springer, 2014.
- [3] Avinava Dubey, Soumen Chakrabarti, and Chiranjib Bhattacharyya. Diversity in ranking via resistive graph centers. In *Proceedings of SIGKDD’11*, pages 78–86. ACM, 2011.

- [4] F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *TIIS*, 5(4):19, 2016.
- [5] Majid Hatami and Saeid Pashazadeh. Improving results and performance of collaborative filtering-based recommender systems using cuckoo optimization algorithm. *IJCA*, 88(16), 2014.
- [6] Jonathan L Herlocker, Joseph A Konstan, Loren G Terveen, and John T Riedl. Evaluating collaborative filtering recommender systems. *TOIS*, 22(1):5–53, 2004.
- [7] Neil Hurley and Mi Zhang. Novelty and diversity in top-n recommendation—analysis and evaluation. *TOIT*, 10(4):14, 2011.
- [8] Ray Jiang, Silvia Chiappa, Tor Lattimore, Andras Agyorgy, and Pushmeet Kohli. Degenerate feedback loops in recommender systems. *arXiv:1902.10730*, 2019.
- [9] Matevž Kunaver and Tomaž Požrl. Diversity in recommender systems—a survey. *Knowledge-Based Systems*, 123:154–162, 2017.
- [10] Amaury L’Huillier, Sylvain Castagnos, and Anne Boyer. The new challenges when modeling context through diversity over time in recommender systems. In *Proceedings of UMAP’16*, pages 341–344. ACM, 2016.
- [11] Rong-Hua Li and Jeffery Xu Yu. Scalable diversified ranking on large graphs. *TKDE*, 25(9):2133–2146, 2012.
- [12] Sean M McNee, John Riedl, and Joseph A Konstan. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI’06*, pages 1097–1101. ACM, 2006.
- [13] Sharad Nandanwar, Aayush Moroney, and M Narasimha Murty. Fusing diversity in recommendations in heterogeneous information networks. In *Proceedings of WSDM’18*, pages 414–422. ACM, 2018.
- [14] Mark Newman. The configuration model. In *Networks*, chapter 12. Oxford university press, 2018.
- [15] Eli Pariser. *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin, 2011.

- [16] Pedro Ramaciotti-Morales, Robin Lamarche-Perrin, Raphael Fournier-S'niehotta, Remy Poulain, Lionel Tabourier, and Fabien Tarissan. Measuring diversity in heterogeneous information networks. *arXiv:2001.01296*, 2020.
- [17] Francesco Ricci, Lior Rokach, and Bracha Shapira. *Recommender systems handbook*. Springer, 2011.
- [18] Chuan Shi, Binbin Hu, Wayne Xin Zhao, and Philip S Yu. Heterogeneous information network embedding for recommendation. *TKDE*, 31(2):357–370, 2018.
- [19] Chuan Shi, Yitong Li, Jiawei Zhang, Yizhou Sun, and Philip S Yu. A survey of heterogeneous information network analysis. *TKDE*, 29(1):17–37, 2016.
- [20] Chuan Shi and Philip S. Yu. *Heterogeneous information network analysis and applications*. Springer, 2017.
- [21] Chuan Shi, Zhiqiang Zhang, Ping Luo, Philip S Yu, Yading Yue, and Bin Wu. Semantic path based personalized recommendation on weighted heterogeneous information networks. In *Proceedings of CIKM'15*, pages 453–462. ACM, 2015.
- [22] Thiago Silveira, Min Zhang, Xiao Lin, Yiqun Liu, and Shaoping Ma. How good your recommender system is? a survey on evaluations in recommendation. *IJMLC*, 10(5):813–831, 2019.
- [23] Andy Stirling. A general framework for analysing diversity in science, technology and society. *J. R. Soc. Interface*, 4(15):707–719, 2007.
- [24] Yizhou Sun, Jiawei Han, Xifeng Yan, Philip S Yu, and Tianyi Wu. Pathsim: Meta path-based top-k similarity search in heterogeneous information networks. *PVLDB*, 4(11):992–1003, 2011.
- [25] Jie Tang, Sen Wu, Jimeng Sun, and Hang Su. Cross-domain collaboration recommendation. In *Proceedings of SIGKDD'12*, pages 1285–1293. ACM, 2012.
- [26] Hanghang Tong, Jingrui He, Zhen Wen, Ravi Konuru, and Ching-Yung Lin. Diversified ranking on large graphs: an optimization viewpoint. In *Proceedings of SIGKDD'11*, pages 1028–1036. ACM, 2011.

- [27] Xiwang Yang, Harald Steck, and Yong Liu. Circle-based recommendation in online social networks. In *Proceedings of SIGKDD'12*, pages 1267–1275. ACM, 2012.
- [28] Xiao Yu, Xiang Ren, Yizhou Sun, Quanquan Gu, Bradley Sturt, Urvashi Khandelwal, Brandon Norick, and Jiawei Han. Personalized entity recommendation: A heterogeneous information network approach. In *Proceedings of WSDM'14*, pages 283–292. ACM, 2014.
- [29] Xiao Yu, Xiang Ren, Yizhou Sun, Bradley Sturt, Urvashi Khandelwal, Quanquan Gu, Brandon Norick, and Jiawei Han. Recommendation in heterogeneous information networks with implicit user feedback. In *Proceedings of RecSys'13*, pages 347–350. ACM, 2013.
- [30] R. Zafarani and H. Liu. Social computing data repository at ASU, 2009.
- [31] Yi Zhang, Jamie Callan, Jamie Callan, and Thomas Minka. Novelty and redundancy detection in adaptive filtering. In *Proceedings of SIGIR'02*, pages 81–88. ACM, 2002.
- [32] Tao Zhou, Zoltán Kuzsók, Jian-Guo Liu, Matúš Medo, Joseph Rushton Wakeling, and Yi-Cheng Zhang. Solving the apparent diversity-accuracy dilemma of recommender systems. *PNAS*, 107(10):4511–4515, 2010.
- [33] Cai-Nicolas Ziegler, Sean M McNee, Joseph A Konstan, and Georg Lausen. Improving recommendation lists through topic diversification. In *Proceedings of WWW'05*, pages 22–32. ACM, 2005.