



AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Tag-based user profiling: A game theoretic approach

This is a pre print version of the following article:		
Original Citation:		
Availability:		
This version is available http://hdl.handle.net/2318/1870191	since 2022-07-20T08:45:26Z	
Publisher:		
Association for Computing Machinery, Inc		
Published version:		
DOI:10.1145/3314183.3323462		
Terms of use:		
Open Access		
Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.		

(Article begins on next page)

Tag-Based User Profiling: A Game Theoretic Approach

Guglielmo Faggioli Mirko Polato Fabio Aiolli ggfaggioli@gmail.com mpolato@math.unipd.it aiolli@math.unipd.it University of Padova Padova, Italy

ABSTRACT

As already pointed out by a constantly growing literature, explainability in recommender systems field is a key aspect to increase users' satisfaction. With the increase of user generated content, tags have proven to be highly relevant when it comes to describe either users or items. A number of strategies that rely on tags have been proposed, yet, many of these algorithms exploit the frequency of user-tags interactions to gain information. We argue that a pure frequentist description might lack of specificity to grasp user's peculiar tastes. Therefore, we propose a novel approach based on game theory that tries to find the best trade-off between generality and detailing. The identified user's description can be used to keep her in the loop and allows the user to have control over system's knowledge. Additionally, we propose a user interface that embeds the proposed user's description and it can be used by the user herself to guide her catalogue's exploration toward novel and serendipitous items.

KEYWORDS

recommender systems, game theory, explanability, clustering, tag, GUI, interface

ACM Reference Format:

Guglielmo Faggioli, Mirko Polato, and Fabio Aiolli. 2019. Tag-Based User Profiling: A Game Theoretic Approach. In *Proceedings of UMAP '19: 27th ACM Conference On User Modelling, Adaptation And Personalization (UMAP '19)*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/nnnnnn. nnnnnnn

1 INTRODUCTION

Likewise search engines a couple of decades ago, Recommender Systems (RSs) are quickly and constantly growing in popularity. We interact with RSs on a daily basis, even when we are not aware of it: e-commerce, streaming services, Ad services, to name a few.

UMAP '19, June 09–12, 2019, Larnaca, Cyprus

 \circledast 2019 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-x-xxxx-x/YY/MM...\$15.00

https://doi.org/10.1145/nnnnnnnnnn

RSs are useful, and often necessary, to guarantee personalized content and to assist users in managing the huge amount of information/items that would be otherwise overwhelming. Pursuing the chimera of a RS able to perfectly model users preferences, a large body of research focused on improving RSs accuracy, i.e., the ability of predicting how much users would enjoy unknown/unseen items. However, despite the initial enthusiasm in developing more and more accurate models, the RS community has started arguing that interacting with a RS is more about the experience as a whole than a mere matter of accurate predictions [14, 16]. This new vision kicked off new research branches focused more on the user satisfaction and less on the *myopic* algorithmic aspect.

Many factors can affect the experience during user interactions with RSs. For example, *transparency* [21, 22] is a key element when it comes to build user trust in the system. Transparency is strictly related to explanability [10, 23], that is, how much a recommendation can be justified in a way that a (preferably, non-expert) user can understand.

Novelty [12] and diversity [11, 26] are also very important to increase user satisfaction. RSs that do not diversify their suggestion tend to create the so-called *filter bubble* [17] (or *portfolio effect* [1]) which can harm the overall user experience.

Involving the user in the recommendation process is an additional factor that can improve the user satisfaction [5]. In particular, providing users with tools to control (or adjust) the recommendation can increase the trust in the system [15]. For example, [2] proposed IntrospectiveViews, a system that provides an interface enabling the user to view and edit her user model. Another example is TasteWeight [3], where the process of recommendation is represented by three connected layers and it enables the user to fine-tune the recommendation. For more comprehensive surveys on user control in RSs please refer to [9] and [13].

In this work, we focus on implicit feedback and we assume that items are described by a set of tags/keywords. We propose a system for improving user experience when interacting with a RSs. The proposed system has two main components: (*i*) a tag-based user modelling, and (*ii*) a web-based GUI that allows users to navigate and (potentially) control the recommendation. The modelling phase is, in turn, composed by two parts: in the first one, tags/keywords are clustered to create *topics* (or meta-tags) which are used to describe the user; the second part learns the user profile. The use of tags for characterizing the user profile is not new in the RS community (e.g., [24]). However, one of the main contributions of our work stands in the way the user profile is learnt. The proposed

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

method, which is based on game theoretical concepts, creates a profile which balances user-wise popular topics, and topics that are instead "rare" but that possibly constitute latent topics the user might not be aware of. This allows users, through the web interface, to explore items that are diverse/novel.

The remainder of the paper is structured as follows. Section 2 describes the method as well as the used dataset. Section 3 describes the web interface and how users can interact with it in order to explore his/her profile. Finally, Section 4 draws some conclusions and suggests some possible future works.

2 METHOD

Our method is based on the concept of tag/keyword. For this reason we assume that an item $i \in I$ is characterized by a set of tags \mathcal{T}_i , and $\mathcal{T} \equiv \bigcup_{i \in I} \mathcal{T}_i$. Tags are human-readable concepts related to items. An item can be described by a number of tags and a tag can be related to more than one item. Due to our case study we will refer to items as movies. Without loss of generality, our approach can be applied to any setting where items are described by tags/keywords. Given a user $u \in \mathcal{U}$, the set of items u has interacted with is indicated by I_u , s.t., $|I_u| = n_u$.

2.1 The dataset

The dataset used in our study, is a merge of two databases about movies: Movielens and TMDb. Movielens has been mainly used to construct the (implicit) rating matrix, while TMDb has been used to extract the tags.

2.1.1 Movielens. Movielens [8] is a freely available dataset often used as the baseline for many recommender algorithms. We used the 20m version, with 20 millions ratings given by approximatively 140k users to 26k movies. Note that, as will be better motivated later, we are not interested in numerical ratings (explicit feedback), but rather in interactions between users and movies (implicit feedback).

2.1.2 *TMDb.* The Movie Database (TMDb¹) is a community built movie and TV database. It represents one of the richest sources of metadata about movies; it contains details of over 26778 movies. Among metadata that can be mined trough TMDb, tags are the most relevant to our task. Tags are generated and applied to movies by users. 17057 tags are available on TMDb. Since many tags are applied to a very small number of movies, we ignored all tags that appear less than 5 times throughout the entire dataset. This left us with 4543 unique tags.

Note that not all movies available in Movielens are also present in TMDb and some movies do not possess any tag, thus being useless in our analysis. Such movies, about 5.5k, have been removed, leaving us with approximatively 21k movies.

2.2 Tag clustering

Usually tags are fine-grained entities that represent highly specific aspects of items, for example, in the movie context a tag can relate to a specific scene. Although very useful, these tags appear to be too specific. There are also many tags that are synonyms or stand for the same concept. Guglielmo Faggioli, Mirko Polato, and Fabio Aiolli

linkage	std. cluster size	min size	max size
complete link	3.28	1	37
avg. link	4.64	1	72
single link	106.70	1	3377

Table 1: Linkage statistics comparison.

For this reason, instead of tags, topics seems more suited for our purposes. Ideally, a topic can be seen as a group of semantically related tags, where each tag represents an aspect of the topic. In order to identify correctly topics related to a specific movie, we developed a clustering strategy over the set of all tags.

The algorithm chosen to build the clustering is hierarchical agglomerative clustering. As suggested by previous work [20], this appears to be the most effective clustering strategy in this context. In the remainder we will refer to the set of all topics (i.e., clusters) with C.

A key step when it comes to employ clustering methodology is to find a good items representation. In this context, the most intuitive way of representing tags consists in the set of movies they belongs to. However, this trivial strategy is unable to correctly dealing with highly popular tags. Moreover, we expect synonyms to be rarely used together: users rarely will suggest new tags that express the very same concept of already available tags, and thus we will rarely observe cooccurrences between synonym words.

In order to account for the above mentioned limitations, we developed a probability-based representation. Specifically, we define $p(t_j|t_i)$ as the probability of observing tag t_j in a movie, given that tag t_i is also present. Note that, if tags t_j and t_k are synonyms, they will appear together with similar tags, although not in the same movies, and thus representations will be similar. Finally, we expect this kind of representation to be more semantically well-founded than pure cooccurrences: two tags are similar if they tend to be used together with the same set of tags. Hence, the representation for tag j is $\tau_j \in \mathbb{R}^{|\mathcal{T}|}$, where $\tau_{j,i} = p(t_j|t_i)$.

The second aspect that needs to be defined, is the similarity measure that is used to compare representations of different entities. Among the plethora of possible similarity measures, we decided to opt for cosine similarity. Cosine similarity appears to be well suited choice considering how we represent tags.

In order to select the linkage strategy we compared different strategies in terms of variance of the cluster sizes. In this way, we wanted to avoid having weakly specialized topics. As shown in Table 1, it appears that the linkage technique able to distribute more evenly tags across topics is *complete link*. The complete linkage produces both clusters where the size has less variance and the smallest largest cluster.

We decided to set the number of topics to 1000 because, empirically, this number leads to clusters that are at the same time specific and general enough.

Table 2 shows some examples of clusters obtained with the just described clustering technique. The *cluster name* is the most popular tag among the ones belonging to the cluster.

¹https://www.themoviedb.org/

cluster tags
british army, british empire, colonialism, dublin, independence, ireland
beguilement, double life, dual identity, envy, false identity, jealousy, new identity, prosecution, rejection, secret
identity
corruption, heist, interrogation, mafia, mobster, organized crime, retired, safecracker
christmas, christmas eve, christmas gift, christmas party, christmas tree, department store, gift, holiday, home
alone, home invasion, little boy, news, broadcast, north pole, precocious child, reindeer, santa claus, sleigh, tele-
caster
artificial intelligence, computer, computer virus, gnosticism, key, man vs machine, mission, rave, super computer,
temple, trans director, truth, underground world, virtual reality
aspiring singer, coming of age, dj, first love, high school, student, modern fairy tale, popularity, principal, promis-
cuity, shyness, summer, summer job, teen comedy, teen drama, teen movie, teenage romance, teenage sexuality,
teenager, virginity
dog, lassie, pet, talking dog
backstage, broadway, broadway musical, broadway show, chorus girl, musical revue, show business, stage play,
stage show, vaudeville

Table 2: Some examples of clusters obtained by our clustering strategy.

2.3 User profiling

Our profiling strategy aims to find a tradeoff between a frequentist and a peculiar description of the user, able to grasp facets of every consumed items. We should consider two aspects of the problem: user profile needs to be (*i*) general enough to describe her, choosing topics that are popular among those she interacted with; (*ii*) highly specific, being able to include in the description topics rarely met, but yet potentially important to the user.

Such profiling approach can be cast into a specific instance of a two-players zero-sum game. Two players want to describe the user: one player chooses topics in order to maximize the number of movies that are covered (frequentist description), while the opponent chooses over movies, selecting those ones that can't be covered by considering only popular topics (movies that can be seen as a niche for the user). This competitive environment, forces the row player to select those topics that are both general, but also able to include specific tastes.

Formally, a strategic form of a two-players zero-sum game is defined by a matrix **M**, called *payoff matrix*.

The game takes place in rounds in which both players, the row player P and the column player Q, play simultaneously. Row player picks a row and column player picks a column of $\mathbf{M} \in \mathbb{R}^{P \times Q}$, where $P = n_u$ and Q = |C|. Each matrix entry $\mathbf{M}_{i,j}$ represents the loss of P, or equivalently the payoff of Q, when the strategies *i* and *j* are played by the two-players. The goal of the player P is to define a strategy that minimizes its expected loss *V*. Conversely, the player Q aims at finding a strategy that maximizes its payoff. Typically, the players play according to mixed strategies: player P selects a row according to a probability distribution \mathbf{p} over the rows, and, similarly, player Q selects a column according to a probability distribution \mathbf{q} over the columns. We refer to the vectors \mathbf{p} and \mathbf{q} as stochastic vectors, that is $\mathbf{p} \in \mathcal{S}_P$ and $\mathbf{q} \in \mathcal{S}_Q$, where $\mathcal{S}_P = \{\mathbf{p} \in \mathbb{R}^P_+ \mid \|\mathbf{p}\|_1 = 1\}$ and $\mathcal{S}_Q = \{\mathbf{q} \in \mathbb{R}^Q_+ \mid \|\mathbf{q}\|_1 = 1\}$.

The optimal pair of strategies $(\mathbf{p}^*,\mathbf{q}^*)$ has a well-know formulation [25], that is

$$V^* = \mathbf{p^*}^\top \mathbf{M} \mathbf{q}^* = \min_{\mathbf{p}} \max_{\mathbf{q}} \mathbf{p}^\top \mathbf{M} \mathbf{q} = \max_{\mathbf{q}} \min_{\mathbf{p}} \mathbf{p}^\top \mathbf{M} \mathbf{q}$$

It is possible, using linear programming, to find the exact distributions **p** and **q**. Given the size of **M**, the exact solution's computation tends to be too expensive and it has shown numerical instability in this task. Thus our approach adopts an approximated solution. In particular, we decided to use the *fictitious play* algorithm [4]. Fictitious play (FP) computes the mixed strategies **p** and **q** incrementally, through a greedy approach. Note that other algorithms might have been used (i.e., [6, 7]), however FP has recently shown performance on big matrices [18].

At the end of FP approximation, we obtain a distribution \mathbf{q} over the set of possible topics, in which element *i* represents the relevance of topic *i* for user. Note that, as already pointed out, these weights, in general, should be different from the frequency of appearance and should be able to describe also those movies that otherwise would have been ignored since they are niche.

2.3.1 The payoff matrix. The payoff matrix **M** is used to describe the relationship between the user and topics that appear in previously seen movies. First of all, it is important to note that we didn't consider only movies the user liked, but rather we considered all movies watched by the user. The reason behind this choice is that we expect the user to interact only with movies that have characteristics and topics that at least the user doesn't dislike: consider an easily impressionable user: we can assume she will never interact with horror or gore movies and thus she won't even give them a rating. It is much more unusual the case where a user that doesn't like horror movies watches them and provide a low rating, simply because she doesn't like the topics.

Given that, each entry *i*, *j* of **M** is defined as follows:

$$\mathbf{M}_{i,j} = \frac{|\mathcal{T}_{w_j} \cap \mathcal{T}_{m_i}|}{|\mathcal{T}_{m_i}|}$$

where \mathcal{T}_{m_i} is the set of tags associated with movie m_i and \mathcal{T}_{w_j} is the set of tags associated with topic w_j . Observe that $||\mathbf{M}_{i,:}||_1 = 1$, $\forall i \in \{1, ..., n_u\}$. We have chosen to use this kind of normalization, due to the fact that, considering pure frequency of topics will lead to biased strategy for the row player that favours undertagged movies.

3 USER INTERFACE

The purpose of the proposed interface is two-fold: first, our interface shows the user what the systems knows about her, second, the interface can be used to guide the user exploring the catalogue that is novelty and diversity driven.



Figure 1: Descriptive area of the user interface.

Figure 1 shows the portion of the interface used to describe the user. On the left hand side of the interface we have a polar chart that is used to show which are the most relevant topics for the user, and in which extent. It is possible to observe that each slice of the chart represents a topic (identified by the most user-wise popular tag belonging to that topic). The radius of the slice indicates the relevance of such topic for the user.

On the right hand side of the "descriptive area" of the interface, it is shown the list of tags that form the specific topic. It is noteworthy that these tags are not the only ones the user interacted with, but all tags related to the topic. This enhances the possibility for the user to explore the catalogue, for example by seeing new tags, still related to topics that are relevant for her, but yet unexplored.

The second part of our interface, as shown in Figure 2, represents the explorative side of our system. The GUI presents a carousel of already seen movies related to the selected topic.

Then, the user will receive a group of recommendations of previously unseen movies that highly relate to those already seen for the selected topic. More so, our interface aims to increase "information scent" and reduce Information Access Cost (IAC), as suggested in [19]. To increase information scent, we show the biggest possible quantity of details for each movie (poster, name, tags - either related to the topic (coloured) or not-) without encumbering the interface. To reduce IAC, a pop-up with other information (genres, synopsis, main cast) is shown by simply hover the movie's area.

Finally, we expect the user to be able to interact with this interface, for example by selecting multiple topics and receive recommendations for movies that are highly related to all selected topics.

4 CONCLUSIONS AND FUTURE WORKS

We presented a system able to describe users via topics (grouping of tags associated with entities the user interacted with). Our

Guglielmo Faggioli, Mirko Polato, and Fabio Aiolli



Figure 2: User interface with both descriptive and explorative/recommendation areas.

system exploits game theoretical concepts in order to find a description that is both general enough, yet capable to gather user's specific facets. Our system embodies many of the state-of-the-art practices to satisfy users. It grants high explainability to a recommender system, showing a precise description of the user. The system exploits important notions, like IAC and information scent, to favour user's exploration. Finally, our system aims to maximize novelty (and possibly serendipity), driving the user toward items that present features the user is known to be interested in, mixed with new and unexpected ones. Although preliminary results suggest a promising research field, we plan to deeply study many different aspects of our methodology. We plan to test other clustering strategies and tags representations, for example we want to observe how the user's description changes when semantic representations are adopted. We plan to expand our dataset, including new sources, like semantic knowledge bases. We plan to evaluate our algorithm in a real setting, with a user study. Finally, we aim to study the generalization capability of our algorithm. At the moment, the trade-off between generality and specificity is fixed, we plan to study a way to make this trade-off parametrical.

REFERENCES

- [1] Kamal Ali and Wijnand van Stam. 2004. TiVo: Making Show Recommendations Using a Distributed Collaborative Filtering Architecture. In Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '04). ACM, New York, NY, USA, 394–401. https://doi.org/10.1145/ 1014052.1014097
- [2] Fedor Bakalov, Birgitta König-Ries, Andreas Nauerz, and Martin Welsch. 2010. IntrospectiveViews: An Interface for Scrutinizing Semantic User Models. In User Modeling, Adaptation, and Personalization, Paul De Bra, Alfred Kobsa, and David Chin (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 219–230.

Tag-Based User Profiling: A Game Theoretic Approach

UMAP '19, June 09-12, 2019, Larnaca, Cyprus

- [3] Svetlin Bostandjiev, John O'Donovan, and Tobias Höllerer. 2012. TasteWeights: A Visual Interactive Hybrid Recommender System. In Proceedings of the Sixth ACM Conference on Recommender Systems (RecSys '12). ACM, New York, NY, USA, 35–42. https://doi.org/10.1145/2365952.2365964
- G.W. Brown. 1951. Iterative solution of games by fictitious play. Activity Analysis of Production and Allocation (1951), 374–376.
- [5] Simon Dooms, Toon De Pessemier, and Luc Martens. 2014. Improving IMDb movie recommendations with interactive settings and filters. In Poster Proceedings of the 8th ACM Conference on Recommender Systems, Vol. 1247. 2.
- [6] Yoav Freund and Robert E. Schapire. 1996. Game Theory, On-line Prediction and Boosting. In Proceedings of the Ninth Annual Conference on Computational Learning Theory (COLT '96). ACM, New York, NY, USA, 325–332. https://doi. org/10.1145/238061.238163
- [7] Yoav Freund and Robert E. Schapire. 1999. Adaptive Game Playing Using Multiplicative Weights. Games and Economic Behavior 29, 1 (1999), 79 – 103. https://doi.org/10.1006/game.1999.0738
- [8] F. Maxwell Harper and Joseph A. Konstan. 2015. The MovieLens Datasets: History and Context. ACM Transactions on Interactive Intelligent Systems 5, 4, Article 19 (2015), 19 pages.
- [9] Chen He, Denis Parra, and Katrien Verbert. 2016. Interactive recommender systems: A survey of the state of the art and future research challenges and opportunities. Expert Syst. Appl. 56 (2016), 9–27.
- [10] Jonathan L. Herlocker, Joseph A. Konstan, and John Riedl. 2000. Explaining Collaborative Filtering Recommendations. In Proceedings of the 2000 ACM Conference on Computer Supported Cooperative Work (CSCW '00). ACM, New York, NY, USA, 241–250. https://doi.org/10.1145/358916.358995
- [11] Jonathan L. Herlocker, Joseph A. Konstan, Loren G. Terveen, and John T. Riedl. 2004. Evaluating Collaborative Filtering Recommender Systems. ACM Transactions on Information Systems 22, 1 (2004), 5–53. https://doi.org/10.1145/963770. 963772
- [12] Rong Hu and Pearl Pu. 2011. Helping Users Perceive Recommendation Diversity. In DiveRS at RecSys.
- [13] Dietmar Jannach, Sidra Naveed, and Michael Jugovac. 2017. User Control in Recommender Systems: Overview and Interaction Challenges. In *E-Commerce* and Web Technologies, Derek Bridge and Heiner Stuckenschmidt (Eds.). Springer International Publishing, Cham, 21–33.
- [14] Joseph A. Konstan and John Riedl. 2012. Recommender systems: from algorithms to user experience. User Modeling and User-Adapted Interaction 22, 1 (01 Apr 2012), 101–123.
- [15] Sean M. McNee, Shyong K. Lam, Joseph A. Konstan, and John Riedl. 2003. Interfaces for Eliciting New User Preferences in Recommender Systems. In Proceedings of the 9th International Conference on User Modeling (UM'03). Springer-Verlag, Berlin, Heidelberg, 178–187.
- [16] Sean M. McNee, John Riedl, and Joseph A. Konstan. 2006. Being Accurate is Not Enough: How Accuracy Metrics Have Hurt Recommender Systems. In CHI '06 Extended Abstracts on Human Factors in Computing Systems (CHI EA '06). ACM, New York, NY, USA, 1097–1101. https://doi.org/10.1145/1125451.1125659
- [17] Eli Pariser. 2011. The Filter Bubble: What the Internet Is Hiding from You. Penguin Group, The.
- [18] Mirko Polato and Fabio Aiolli. 2018. Interpretable preference learning: a game theoretic framework for large margin on-line feature and rule learning. In AAAI'19.
- [19] Tobias Schnabel, Paul Bennett, and Thorsten Joachims. 2018. Improving Recommender Systems Beyond the Algorithm. (02 2018).
- [20] Andriy Shepitsen, Jonathan Gemmell, Bamshad Mobasher, and Robin Burke. 2008. Personalized Recommendation in Social Tagging Systems Using Hierarchical Clustering. In Proceedings of the 2008 ACM Conference on Recommender Systems (RecSys '08). ACM, New York, NY, USA, 259–266. https://doi.org/10. 1145/1454008.1454048
- [21] K. Swearingen and R. Sinha. 2001. Beyond Algorithms: An HCI Perspective on Recommender Systems. In ACM SIGIR. Workshop on Recommender Systems, Vol. Vol. 13, Numbers 5-6. 393–408.
- [22] Nava Tintarev and Judith Masthoff. 2007. Effective Explanations of Recommendations: User-centered Design. In Proceedings of the 2007 ACM Conference on Recommender Systems (RecSys '07). ACM, New York, NY, USA, 153–156. https://doi.org/10.1145/1297231.1297259
- [23] Nava Tintarev and Judith Masthoff. 2007. A Survey of Explanations in Recommender Systems. In Proceedings of the 2007 IEEE 23rd International Conference on Data Engineering Workshop (ICDEW '07). IEEE Computer Society, Washington, DC, USA, 801–810. https://doi.org/10.1109/ICDEW.2007.4401070
- [24] Jesse Vig, Shilad Sen, and John Riedl. 2009. Tagsplanations: Explaining Recommendations Using Tags. In Proceedings of the 14th International Conference on Intelligent User Interfaces (IUI '09). ACM, New York, NY, USA, 47–56. https: //doi.org/10.1145/1502650.1502661
- [25] John von Neumann. 1928. Zur Theorie der Gesellschaftsspiele. 100 (1928), 295– 320.
- [26] Cai-Nicolas Ziegler, Sean M. McNee, Joseph A. Konstan, and Georg Lausen. 2005. Improving Recommendation Lists Through Topic Diversification. In Proceedings

of the 14th International Conference on World Wide Web (WWW '05). ACM, New York, NY, USA, 22–32. https://doi.org/10.1145/1060745.1060754