

# Towards Evaluating User Profiling Methods Based on Explicit Ratings on Item Features

Luca Luciano Costanzo  
Politecnico di Milano, Italy  
lucaluciano.costanzo@mail.polimi.it

Yashar Deldjoo  
Polytechnic University of Bari, Italy  
deldjooy@acm.org

Maurizio Ferrari Dacrema  
Politecnico di Milano, Italy  
maurizio.ferrari@polimi.it

Markus Schedl  
Johannes Kepler University Linz,  
Austria  
markus.schedl@jku.at

Paolo Cremonesi  
Politecnico di Milano, Italy  
paolo.cremonesi@polimi.it

## ABSTRACT

In order to improve the accuracy of recommendations, many recommender systems nowadays use side information beyond the user rating matrix, such as item content. These systems build user profiles as estimates of users' interest on content (e.g., movie genre, director or cast) and then evaluate the performance of the recommender system as a whole e.g., by their ability to recommend relevant and novel items to the target user. The user profile modelling stage, which is a key stage in content-driven RS is barely properly evaluated due to the lack of publicly available datasets that contain user preferences on content features of items.

To raise awareness of this fact, we investigate differences between *explicit* user preferences and *implicit* user profiles. We create a dataset of explicit preferences towards content features of movies, which we release publicly. We then compare the collected explicit user feature preferences and implicit user profiles built via state-of-the-art user profiling models. Our results show a maximum average pairwise cosine similarity of 58.07% between the explicit feature preferences and the implicit user profiles modelled by the best investigated profiling method and considering movies' genres only. For actors and directors, this maximum similarity is only 9.13% and 17.24%, respectively. This low similarity between explicit and implicit preference models encourages a more in-depth study to investigate and improve this important user profile modelling step, which will eventually translate into better recommendations.

## ACM Reference Format:

Luca Luciano Costanzo, Yashar Deldjoo, Maurizio Ferrari Dacrema, Markus Schedl, and Paolo Cremonesi. 2019. Towards Evaluating User Profiling Methods Based on Explicit Ratings on Item Features. In *Proceedings of Joint Workshop on Interfaces and Human Decision Making for Recommender Systems (IntRS '19)*. CEUR-WS.org, 5 pages.

## 1 INTRODUCTION

The performance of collaborative filtering (CF) recommendation models have reached a remarkable level of maturity. These models are now widely adopted in real-world recommendation engines because of their state-of-the-art recommendation quality. In recent years, a number of recommendation scenarios have emerged, which have encouraged the research community to consider using various

additional information sources (aka side information) beyond the user rating matrix [25]. A prominent example—and the one we focus on—is item content. In the movie domain, for instance, a variety of content features have been considered, such as metadata or features extracted directly from the core audio-visual signals. Metadata-based movie recommender systems typically use genre [10, 13, 26] or user-generated tags [18, 31, 34] over which user profiles are built, assuming that these aspects represent the semantic content of movies. In contrast, audio-visual signals represent the low-level content (e.g., color, lighting, spoken dialogues, music, etc.) [4, 6, 7, 9, 10]. Some approaches try to infer semantic concepts from low-level representations, e.g., via word2vec embeddings [2], deep neural networks [30, 33], fuzzy logic [28], or genetic algorithms [20]. For these reasons, it is evident that item content plays a key role in building hybrid or content-based filtering (CBF) models and, furthermore, it is important to correctly distinguish and weight the item features by their estimated relevance for a target user, to better model his or her tastes.

In Figure 1, we illustrate a simplified diagram that shows our research contributions. Standard recommendation based on content (CBF or hybrid) is structured in three main steps: (i) extraction of item content, consisting of building a *feature vector* that describes each item  $i$ ; (ii) building the *profile of the target user*  $\mathbf{p}_u$ , i.e., a structured representation of the user's preference over item content features; (iii) matching the user profile  $\mathbf{p}_u$  against the feature vector of each item  $\mathbf{f}_i$  to produce the list of recommended items most similar to the target user's tastes.

A shortcoming of typical RS evaluation is that the user profiling stage, which is a key part of the RS, is barely evaluated. Usually, only the performance of the entire RS, which is composed of several components, is assessed and how effectively the user profiling step functions remains an open question. We argue that it is important to investigate the user profiling stage and compare performance of different profile modelling methods (see upper part of Figure 1).

The goal of this work is therefore to investigate the difference between explicit user ratings on individual movie content features (e.g., genre, actors, or directors) and implicit models inferred via state-of-the-art user modelling techniques from explicit ratings of the whole movies. To this end, we (i) create (and make publicly available) a varied dataset of explicit ratings both on movies and content features and (ii) evaluate different user profiling methods and compare their resulting implicit models against the true feature ratings provided in the collected dataset.

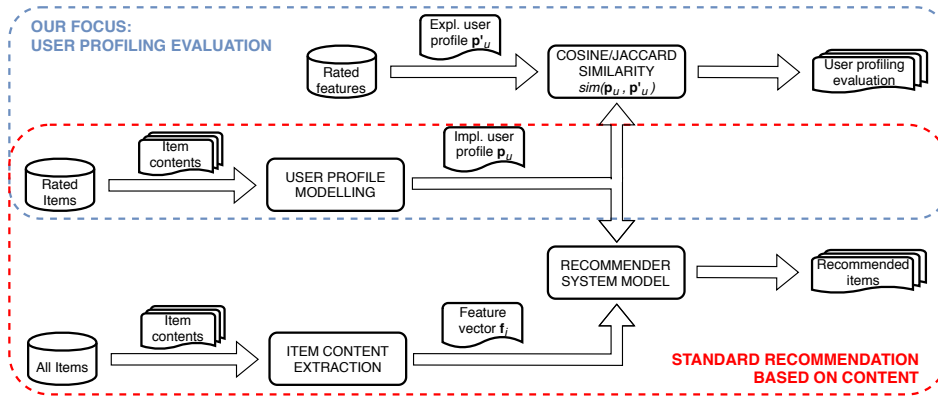


Figure 1: Main steps involved in a recommendation system leveraging content information, highlighting our contributions.

## 2 RELATED WORK

With respect to previous research, to the best of our knowledge, the only work that evaluates implicit user profiles against true ratings on content features is [21]. Nasery et al. compare actually rated features with the ones implicitly derived from rated movies, but no concrete user profiling methods are investigated. Instead, the number of times each feature is explicitly rated and the number of times it appears in the content of all rated movies is counted, and these counts are compared. The authors create a dataset of movies’ feature ratings (genres, actors/cast, and directors), dubbed PoliMovie,<sup>1</sup> through a survey web application they built. Their approach, using limited survey questions and a fixed reduced dataset of top popular movies and features, extracted from IMDb,<sup>2</sup> tends to push users to limited and convergent preferences. In contrast, we systematically investigate 4 methods to model implicit user profiles and we compare them with explicit user profiles obtained by feature ratings. Another contribution of the work at hand is the creation of a dataset that includes ratings on movie content features. Other datasets commonly used in movie recommender systems research, but which do not contain such feature ratings, include MovieLens 20M (ML-20M) [11], IMDB Movies Dataset [16], The Movies Dataset [3], MMTF-14K and MVCD-7K [5, 8] and the Netflix Prize dataset [22].

## 3 USER PROFILE MODELLING TECHNIQUES

To create a *user profile*, we adopt the *vector profile* representation, consisting of weighted attributes measuring the user’s taste on each feature [6, 14], because it is best suited for our evaluation in terms of similarity functions. Formally, the user profiling methods we investigate build the user profile  $\mathbf{p}_u$  as a vector whose attributes are the relevance weight of each feature  $f$  for the target user  $u$ , denoted as  $h_{u,f}$ .

We analyze 3 state-of-the-art methods from literature to model user profiles and we refer to them according to the first author of the corresponding publication, for simplicity and a 4<sup>th</sup> method that applies the TF-IDF (term frequency–inverse document frequency) term weighting idea, which is widely used in CBF and, in general, in information retrieval [19, 23, 29].

**Zhang method.** Zhang et al. [32] build the user profile based on item ratings or explicit feature ratings. Let  $\mathcal{U}$  and  $\mathcal{I}$  denote the set of users and items, respectively, and  $\mathcal{F}$  the set of all features of the items. In case of binary ratings (like in our dataset), this method assigns relevance weight  $h_{u,f}$  equal to 1 for each feature  $f$  in  $\mathcal{F}$  that applies to items with which the target user  $u$  interacted with, 0 otherwise. The obvious limitation of this method is that it assigns only weights 0 or 1 to the features, without distinguishing their relevance for the user.

**Li method.** Li et al. [17], unlike Zhang et al., differentiate the relevance of features contained in an item by assigning scalar weights. Their method furthermore ignores items with low ratings by using a threshold value. In case of binary ratings, the threshold rating  $r_\tau$  is 0 and the relevance weight  $h_{u,f}$  of each feature  $f$  in  $\mathcal{F}$  for the target user  $u$  becomes the percentage of occurrences of  $f$  in the items  $u$  interacted with:  $h_{u,f} = N_{u,f}/M_u$ , where  $N_{u,f}$  is the number of items rated by user  $u$  containing feature  $f$  and  $M_u$  is the total number of items rated by user  $u$ .

**Symeonidis method.** Symeonidis et al. [27] adopt an approach similar to TF-IDF to compute feature relevance weights, but define them in the vector space of user profiles. The rationale of using TF-IDF is to increase the relevance of rare features contained in less user profiles. Symeonidis et al. also use a fixed rating threshold to consider only the most relevant items. In case of binary ratings, the threshold rating  $r_\tau$  is set to 0 and the relevance weight  $h_{u,f}$  of each feature  $f$  in  $\mathcal{F}$  for the target user  $u$  is computed as:  $h_{u,i} = FF(u, f) \cdot IUF(f)$ , where  $FF(u, f)$  is the feature frequency, i.e., the number of times feature  $f$  occurs in movies rated by  $u$ , and  $IUF(f)$  is the inverse user frequency of feature  $f$ .  $IUF(f) = \log \frac{|\mathcal{U}|}{UF(f)}$ , where  $UF(f)$  is the user frequency of  $f$ , i.e., the number of users whose rated movies contain feature  $f$  at least once.

**TF-IDF method.** After having reviewed the 3 state-of-art methods described above, we decided to investigate another variant of TF-IDF as a user profiling method. The Symeonidis method above is similar to TF-IDF, but it is user-centric because it considers the vector space of user profiles. Instead, our proposed TF-IDF method is item-centric as it considers the vector space of items (movies). First, we compute the IDF of each feature  $f$  as:  $IDF(f) = \log \frac{|\mathcal{I}|}{n_f}$ , where  $n_f$  denotes the number of items in  $\mathcal{I}$  in which feature  $f$  occurs at least once. Then, for each user  $u$ , we compute the relevance weight

<sup>1</sup>PoliMovie: <http://bit.ly/polimovie>

<sup>2</sup>Internet Movie Database (IMDB): [www.imdb.com](http://www.imdb.com)

$h_{u,f}$  of a feature  $f$  as:  $h_{u,f} = TF(u, f) \cdot IDF(f)$ , where  $TF(u, f)$  is equivalent to  $FF(u, f)$  of the Symeonidis method (i.e., number of times feature  $f$  occurs in items rated by user  $u$ ). In contrast to the method by Symeonidis et al.,  $IDF(f)$  is computed in relation to all the existing items in which feature  $f$  appears, not related to user profiles. As will be shown in Section 5.2, our TF-IDF method yields better results than Symeonidis et al.’s.

## 4 DATA ACQUISITION

The dataset we use to evaluate user profiling methods has been collected through a web application we implemented, which can be navigated on a variety of stationary and mobile devices. It provides access to a large catalogue of more than 450K movies and any related content feature. This vast breadth of choice is possible thanks to the fact that we retrieve up-to-date information on-the-fly from TMDb<sup>3</sup> via APIs. We developed the application with the idea of a completely free user experience, instead of making it like a survey application, so that users are not forced in any way during their selections.

To acquire the needed data, we asked users to select a set of “favourites”, which included at least 5 movies, 2 genres, 3 actors, and 1 director. Users were, nevertheless, free to select more than these numbers of elements. We also asked users to provide some demographics information: age range, gender, and country of residence. The collection of data was divided into two phases, the first one involved the *volunteer* users, which are the ones invited to freely contribute (friends, family, acquaintances, and colleagues of the authors), while the second phase involved users recruited by the *crowdsourcing* platform MTurk,<sup>4</sup> which have been paid between 20 and 50 US cents for their contribution. To assess the participants’ reliability, we also asked them to complete a final consistency test which required to select again all (and only) the favourites they remember to have added (from a list of movies, genres, and actors of random popular elements). A user’s reliability is then estimated by means of the precision score computed on the re-selection of correct favourites.

Finally, in order to explore a catalogue of existing features needed for user profiling evaluation, we retrieved The Movies Dataset [3] containing the content of 45,3K movies scraped from TMDb. Then, we extended this dataset by scraping the content of missing movies that were added as favorites by users on our web application.

**Dataset characteristics.** We have collected the preferences of 194 users, 180 (93%) of whom have added the minimum number of required favourites. Among all users, 81 (42%) are volunteers and 113 (58%) are paid ones. We consider users reliable if they are either volunteers that have completed the required favourites or crowdsourced users who scored at least 50% of precision during the consistency test (see above). The reliable volunteers are 67 (83% of all volunteers), while the crowdsourcing ones are 88 (78% of all crowdsourcing), hence a total of 155 reliable users (80% of all users).

Regarding users’ gender, 115 users (59%) are male, 66 are female (34%), and 13 (7%) did not specify gender. 53% of the users are between 24 and 30 years old. We received registrations from users coming from 10 different countries, mainly from Italy (40%), India (31%), and United States (19%).

<sup>3</sup>The Movie Database (TMDb): [www.themoviedb.org](http://www.themoviedb.org)

<sup>4</sup>Amazon Mechanical Turk (MTurk): [www.mturk.com](http://www.mturk.com)

We collected a total 4,109 favourites (movies and content features) selected by participants, including 1,212 unique elements, i.e., favourites selected by at least one user. In the following experiments, we include only favourites of *reliable* users, that are 3,341 (81%), including 1,737 favourite movies, 461 genres, 698 actors, 198 directors, 74 production companies, 92 production countries, 39 producers, 17 screenwriters, 21 release years, and 4 sound crew members. The dataset is available on Kaggle<sup>5</sup>.

## 5 RESULTS AND DISCUSSION

### 5.1 Initial statistical analysis

An initial statistical analysis highlights main differences between the set of all explicitly rated features and the set of all implicit features extracted from rated movies. In Tables 1 and 2, we present a comparison between the explicit and implicit sets of features, in percentage of common attributes (features), focusing on the  $k$  most frequently selected attributes, respectively, for genre, actor, and director. These tables generally highlight a low overlap between the explicitly preferred features and the implicitly estimated ones (derived from favourite movies), in particular for actors and directors. The only exception is the genre attribute, which reveals a maximum overlap of 94.74% when considering all 19 genres. These results generally confirm the previous findings in [21] regarding existing gaps between explicitly selected features and implicitly estimated ones, with a different dataset containing more up-to-date movies and not limited to the most popular movies as used in [21].

**Table 1: Common genres in the most selected  $k$  attributes, either explicitly or implicitly**

$k$	No. common genres	% common genres
5	3	60.00%
10	8	80.00%
15	13	86.67%
All genres (19)	18	94.74%

We further provide a finer-grained analysis of the gap between explicit and implicit preferences of users according to their gender. In Tables 3 and 4, we compare the 5 most frequently selected genres, actors, and directors, by *male* and *female* users, respectively. We notice a substantial difference between between male and female users with the exception of genre.

Investigating the results, it is surprising that in both Tables 3 and 4 Stan Lee is among the top *implicitly* preferred actors even if he barely acted as a main character in any movie. The most probable reason is that even though he has not been selected explicitly as favourite actor by study participants, he appeared in all Marvel movies (in small “cameo roles”), so he is included in the implicit profiles. Furthermore, it is surprising that the genre “action” is highly ranked by female users. This could be due to the fact that the genre tastes of young women might be changing nowadays, especially because many popular action movies, like the Marvel ones, are liked by many people (especially under 30, i.e., the largest age group in our dataset), irrespective of gender. Nonetheless, the other differences between male and female users suggest to embed gender information in a recommender system.

<sup>5</sup><https://www.kaggle.com/lucacostanzo/mints-dataset-for-recommender-systems>

**Table 2: Common quota of either actors or directors between the most selected  $k$  attributes, either explicitly or implicitly**

$k$	% of common actors	% of common directors
10	10.00%	30.00%
20	20.00%	20.00%
40	22.50%	27.50%
60	16.67%	28.33%

**Table 3: Most selected 5 features, either explicitly ( $R_f^{exp}$ ) or implicitly ( $R_f^{imp}$ ), by male users;**

	Pos.	Explicit selection	$R_f^{exp}$	Implicit selection	$R_f^{imp}$
Genres	1	Action	51	Action	86
	2	Drama	31	Adventure	83
	3	Adventure	30	Drama	80
	4	Thriller	28	Science Fiction	76
	5	Science Fiction	28	Thriller	74
Actors	1	Robert Downey Jr.	16	Samuel L. Jackson	64
	2	Johnny Depp	15	Stan Lee	56
	3	Jason Statham	10	Bradley Cooper	51
	4	Leonardo Di-Caprio	10	Paul Bettany	47
	5	Tom Hardy	8	Vin Diesel	47
Directors	1	Quentin Tarantino	11	Hajar Mainl	42
	2	Steven Spielberg	9	Chris Castaldi	41
	3	Joe Russo	7	Mark Rossini	41
	4	M. Night Shyamalan	6	Lori Grabowski	41
	5	Christopher Nolan	6	Eli Sasich	41

**Table 4: Most selected 5 features, either explicitly ( $R_f^{exp}$ ) or implicitly ( $R_f^{imp}$ ), by female users;**

	Pos.	Explicit selection	$R_f^{exp}$	Implicit selection	$R_f^{imp}$
Genres	1	Drama	26	Drama	52
	2	Action	22	Adventure	48
	3	Adventure	14	Action	47
	4	Comedy	14	Fantasy	45
	5	Thriller	13	Science Fiction	43
Actors	1	Robert Downey Jr.	12	Stan Lee	27
	2	Leonardo Di-Caprio	7	Samuel L. Jackson	26
	3	Jennifer Lawrence	5	Bradley Cooper	23
	4	Chris Hemsworth	5	Djimon Hounsou	21
	5	Bruce Willis	4	James McAvoy	21
Directors	1	Joe Russo	4	Anthony Russo	16
	2	Christopher Nolan	4	Joe Russo	16
	3	Steven Spielberg	4	Bryan Singer	15
	4	Martin Scorsese	2	Hajar Mainl	14
	5	Ridley Scott	2	Chris Castaldi	14

## 5.2 Evaluation of user profiling methods

We study the user profiling step in-depth by investigating the 4 user profiling methods described in Section 3. Our aim is to analyze the similarity (i.e., the overlap) between the implicitly modelled user profiles and the real explicit tastes of users. For each target user  $u$ , we built his or her *explicit profile*  $p'_u$  as vector composed of relevance weights equal to 1, for all the features explicitly rated by  $u$ , and weight 0 for the ones not rated. Then we computed the pairwise similarity between the explicit user profiles and *implicit profiles*  $p_u$  produced by each method, using cosine similarity and Jaccard similarity. The highest is this similarity, the most accurate is the implicit user profile modelled.

The average pairwise similarity  $sim(p_u, p'_u)$  between implicit user profile  $p_u$  and explicit one  $p'_u$  is shown in Table 5. As revealed in the table and already anticipated in Section 3, the TF-IDF method

**Table 5: Average pairwise similarity between explicit and implicit user profiles, for all the methods.**

Similarity	Feature	Zhang	Li	Symeonidis	TF-IDF
Cosine	Genre	48.52%	58.07%	42.00%	53.08%
	Actor	7.03%	9.13%	6.50%	7.24%
	Director	15.17%	17.24%	15.32%	16.14%
Jaccard	Genre	27.49%	36.19%	18.54%	33.36%
	Actor	0.97%	5.73%	2.87%	4.64%
	Director	5.22%	10.24%	6.30%	8.17%

yields better results than Symeonidis even if they are intrinsically similar, hence the item-centric TF-IDF approach outperforms the user profile-based one. In general, the average pairwise similarities are remarkably low, even for the best investigated method, i.e., Li. The overlap between explicit and implicit profiles increases if we consider only genres; the reason is that the catalogue of all possible genres in the dataset is rather limited (19) compared to actors (567K) and directors (58K). The Jaccard measure yields lower similarities because it can be applied only to vectors composed of binary attributes while our tested profiling methods compute scalar weights (except for Zhang); hence we had to cut-off some feature weights by considering only the  $k$  most relevant features in the implicit profile of each user considered, in which  $k$  is the number of explicit features rated by that user.

The presented results underline the low effectiveness of the investigated user profiling methods to model real user tastes. This finding gives rise to the need of further research on this important user profiling step when devising recommender systems. If user profiles are not properly modelled before applying any RS technique, the accuracy of the final recommendations will likely be affected and lowered by an inaccurate representation of the user's tastes.

## 6 CONCLUSION AND FUTURE WORKS

In this paper, we analyzed the user profiling modelling by studying the differences between explicit user preferences and implicit user profiles. We evaluated different user profiling methods and showed that even the best profiling method that we tested provided low pairwise similarities between explicit and implicit profiles. This finding can be explained by the fact that when a user rates a movie, he is implicitly rating only some characteristics of the item that impacted on her (but not all). Also, it could happen that a user may select a movie but she only loved some part of it (e.g., very good director but bad actors), and this can result in the introduction of some noise in the learning process. Overall, our study encourages a more in-depth on ways we can obtain reliable feedbacks on features and study the optimization of the user profile modelling step in RS, which will eventually allow to produce more accurate recommendations. Furthermore, we publicly provide the dataset that we collected and used for evaluation, which includes ratings on movies and on corresponding content features.

In the future, we plan to investigate the generalizability of findings in this work on other domains where the exist a wide variety of item content features and personalization on these features is paramount, in domains including but not limited to fashion [12], music domain [24], tourism [1, 15] and so forth.

## REFERENCES

- [1] Jens Adamczak, Gerard-Paul Leyson, Peter Knees, Yashar Deldjoo, Farshad Bakhshandegan Moghaddam, Julia Neidhardt, Wolfgang Wörndl, and Philipp Monreal. 2019. Session-Based Hotel Recommendations: Challenges and Future Directions. *arXiv preprint arXiv:1908.00071* (2019).
- [2] Fahad Anwaar, Naima Iltaf, Hammad Afzal, and Raheel Nawaz. 2018. HRS-CE: A hybrid framework to integrate content embeddings in recommender systems for cold start items. *J. Comput. Science* 29 (2018), 9–18. DOI: <http://dx.doi.org/10.1016/j.jocs.2018.09.008>
- [3] Rounak Banik. 2017. The Movies Dataset. Dataset on Kaggle. (2017). <https://www.kaggle.com/rounakbanik/the-movies-dataset>
- [4] Yashar Deldjoo, Mihai Gabriel Constantin, Hamid Eghbal-Zadeh, Bogdan Ionescu, Markus Schedl, and Paolo Cremonesi. 2018. Audio-visual encoding of multimedia content for enhancing movie recommendations. In *Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2-7, 2018*, Sole Pera, Michael D. Ekstrand, Xavier Amatriain, and John O'Donovan (Eds.). ACM, 455–459. DOI: <http://dx.doi.org/10.1145/3240323.3240407>
- [5] Yashar Deldjoo, Mihai Gabriel Constantin, Bogdan Ionescu, Markus Schedl, and Paolo Cremonesi. 2018. MMTF-14K: a multifaceted movie trailer feature dataset for recommendation and retrieval. In *Proceedings of the 9th ACM Multimedia Systems Conference*. 450–455. DOI: <http://dx.doi.org/10.1145/3204949.3208141>
- [6] Yashar Deldjoo, Paolo Cremonesi, Markus Schedl, and Massimo Quadrona. 2017. The effect of different video summarization models on the quality of video recommendation based on low-level visual features. In *Proceedings of the 15th International Workshop on Content-Based Multimedia Indexing, CBMI 2017, Florence, Italy, June 19-21, 2017*. ACM, 20:1–20:6. DOI: <http://dx.doi.org/10.1145/3095713.3095734>
- [7] Yashar Deldjoo, Maurizio Ferrari Dacrema, Mihai Gabriel Constantin, Hamid Eghbal-zadeh, Stefano Cereda, Markus Schedl, Bogdan Ionescu, and Paolo Cremonesi. 2019. Movie genome: alleviating new item cold start in movie recommendation. *User Model. User-Adapt. Interact.* 29, 2 (2019), 291–343. DOI: <http://dx.doi.org/10.1007/s11257-019-09221-y>
- [8] Yashar Deldjoo and Markus Schedl. 2019. Retrieving Relevant and Diverse Movie Clips Using the MFVCD-7K Multifaceted Video Clip Dataset. In *Proceedings of the 17th Int. Workshop on Content-Based Multimedia Indexing*.
- [9] Fuhu Deng, Panlong Ren, Zhen Qin, Gu Huang, and Zhiguang Qin. 2018. Leveraging Image Visual Features in Content-Based Recommender System. *Scientific Programming* 2018 (2018), 5497070:1–5497070:8. DOI: <http://dx.doi.org/10.1155/2018/5497070>
- [10] Ralph Jose Rassweiler Filho, Jonatas Wehrmann, and Rodrigo C. Barros. 2017. Leveraging deep visual features for content-based movie recommender systems. In *2017 International Joint Conference on Neural Networks, IJCNN 2017, Anchorage, AK, USA, May 14-19, 2017*. IEEE, 604–611. DOI: <http://dx.doi.org/10.1109/IJCNN.2017.7965908>
- [11] F. Maxwell Harper and Joseph A. Konstan. 2016. The MovieLens Datasets: History and Context. *TiS* 5, 4 (2016), 19:1–19:19. DOI: <http://dx.doi.org/10.1145/2827872>
- [12] Ruining He and Julian McAuley. 2016. Ups and downs: Modeling the visual evolution of fashion trends with one-class collaborative filtering. In *proceedings of the 25th international conference on world wide web*. International World Wide Web Conferences Steering Committee, 507–517.
- [13] Tae-Gyu Hwang, Chan-Soo Park, Jeong-Hwa Hong, and Sung Kwon Kim. 2016. An algorithm for movie classification and recommendation using genre correlation. *Multimedia Tools Appl.* 75, 20 (2016), 12843–12858. DOI: <http://dx.doi.org/10.1007/s11042-016-3526-8>
- [14] Ameni Kacem. 2017. *Personalized Information Retrieval based on Time-Sensitive User Profile. (Recherche d'Information Personnalisée basée sur un Profil Utilisateur Sensible au Temps)*. Ph.D. Dissertation. Paul Sabatier University, Toulouse, France. <https://tel.archives-ouvertes.fr/tel-01707423>
- [15] Peter Knees, Yashar Deldjoo, Farshad Bakhshandegan Moghaddam, Jens Adamczak, Gerard-Paul Leyson, and Philipp Monreal. 2019. RecSys Challenge 2019: Session-based Hotel Recommendations. In *Proceedings of the Thirteenth ACM Conference on Recommender Systems (RecSys '19)*. ACM, New York, NY, USA, 2. DOI: <http://dx.doi.org/10.1145/3298689.3346974>
- [16] Orges Leka. 2016. IMDB Movies Dataset. Dataset on Kaggle. (2016). <https://www.kaggle.com/orgesleka/imdbmovies>
- [17] Qing Li and Byeong Man Kim. 2004. Constructing User Profiles for Collaborative Recommender System. In *Advanced Web Technologies and Applications, 6th Asia-Pacific Web Conference, APWeb 2004, Hangzhou, China, April 14-17, 2004, Proceedings (Lecture Notes in Computer Science)*, Jeffrey Xu Yu, Xuemin Lin, Hongjun Lu, and Yanchun Zhang (Eds.), Vol. 3007. Springer, 100–110. DOI: [http://dx.doi.org/10.1007/978-3-540-24655-8\\_11](http://dx.doi.org/10.1007/978-3-540-24655-8_11)
- [18] Haibo Liu, Shi Feng, and Ge Yu. 2017. An interest propagation based movie recommendation method for social tagging system. In *2017 International Conference on Machine Learning and Cybernetics, ICMLC 2017, Ningbo, China, July 9-12, 2017*. IEEE, 130–135. DOI: <http://dx.doi.org/10.1109/ICMLC.2017.8107754>
- [19] Pasquale Lops, Marco de Gemmis, and Giovanni Semeraro. 2011. Content-based Recommender Systems: State of the Art and Trends. In *Recommender Systems Handbook*, Francesco Ricci, Lior Rokach, Bracha Shapira, and Paul B. Kantor (Eds.). Springer, 73–105. DOI: [http://dx.doi.org/10.1007/978-0-387-85820-3\\_3](http://dx.doi.org/10.1007/978-0-387-85820-3_3)
- [20] Juergen Mueller. 2017. Combining aspects of genetic algorithms with weighted recommender hybridization. In *Proceedings of the 19th International Conference on Information Integration and Web-based Applications & Services, iiWAS 2017, Salzburg, Austria, December 4-6, 2017*, Maria Indrawan-Santiago, Matthias Steinbauer, Ivan Luiz Salvadori, Ismail Khalil, and Gabriele Anderst-Kotsis (Eds.). ACM, 13–22. DOI: <http://dx.doi.org/10.1145/3151759.3151765>
- [21] Mona Nasery, Mehdi Elahi, and Paolo Cremonesi. 2015. PoliMovie: a feature-based dataset for recommender systems. DOI: <http://dx.doi.org/10.13140/RG.2.2.20636.49286>
- [22] Netflix. 2009. Netflix Prize Data. Dataset on Kaggle. (2009). <https://www.kaggle.com/netflix-inc/netflix-prize-data>
- [23] Diego Sánchez-Moreno, María N. Moreno García, Nasim Sonboli, Bamshad Mobasher, and Robin Burke. 2018. Inferring User Expertise from Social Tagging in Music Recommender Systems for Streaming Services. In *Hybrid Artificial Intelligent Systems - 13th International Conference, HAIS 2018, Oviedo, Spain, June 20-22, 2018, Proceedings (Lecture Notes in Computer Science)*, Francisco Javier de Cos Juez, José Ramón Villar, Enrique A. de la Cal, Álvaro Herrero, Héctor Quintián, José Antonio Sáez, and Emilio Corchado (Eds.), Vol. 10870. Springer, 39–49. DOI: [http://dx.doi.org/10.1007/978-3-319-92639-1\\_4](http://dx.doi.org/10.1007/978-3-319-92639-1_4)
- [24] Markus Schedl, Hamed Zamani, Ching-Wei Chen, Yashar Deldjoo, and Mehdi Elahi. 2018. Current challenges and visions in music recommender systems research. *International Journal of Multimedia Information Retrieval* 7, 2 (2018), 95–116.
- [25] Yue Shi, Martha Larson, and Alan Hanjalic. 2014. Collaborative filtering beyond the user-item matrix: A survey of the state of the art and future challenges. *ACM Computing Surveys (CSUR)* 47, 1 (2014), 3.
- [26] Márcio Soares and Paula Viana. 2017. The Semantics of Movie Metadata: Enhancing User Profiling for Hybrid Recommendation. In *Recent Advances in Information Systems and Technologies - Volume 1 [WorldCIST'17, Porto Santo Island, Madeira, Portugal, April 11-13, 2017]*. (Advances in Intelligent Systems and Computing), Álvaro Rocha, Ana Maria R. Correia, Hojjat Adeli, Luís Paulo Reis, and Sandra Costanzo (Eds.), Vol. 569. Springer, 328–338. DOI: [http://dx.doi.org/10.1007/978-3-319-56535-4\\_33](http://dx.doi.org/10.1007/978-3-319-56535-4_33)
- [27] Panagiotis Symeonidis, Alexandros Nanopoulos, and Yannis Manolopoulos. 2007. Feature-Weighted User Model for Recommender Systems. In *User Modeling 2007, 11th International Conference, UM 2007, Corfu, Greece, June 25-29, 2007, Proceedings (Lecture Notes in Computer Science)*, Cristina Conati, Kathleen F. McCoy, and Georgios Paliouras (Eds.), Vol. 4511. Springer, 97–106. DOI: [http://dx.doi.org/10.1007/978-3-540-73078-1\\_13](http://dx.doi.org/10.1007/978-3-540-73078-1_13)
- [28] Pooja Bhatt Vashisth, Purnima Khurana, and Punam Bedi. 2017. A fuzzy hybrid recommender system. *Journal of Intelligent and Fuzzy Systems* 32, 6 (2017), 3945–3960. DOI: <http://dx.doi.org/10.3233/JIFS-14538>
- [29] Donghui Wang, Yanchun Liang, Dong Xu, Xiaoyue Feng, and Renchu Guan. 2018. A content-based recommender system for computer science publications. *Knowl.-Based Syst.* 157 (2018), 1–9. DOI: <http://dx.doi.org/10.1016/j.knsys.2018.05.001>
- [30] Jian Wei, Jianhua He, Kai Chen, Yi Zhou, and Zuoyn Tang. 2017. Collaborative filtering and deep learning based recommendation system for cold start items. *Expert Syst. Appl.* 69 (2017), 29–39. DOI: <http://dx.doi.org/10.1016/j.eswa.2016.09.040>
- [31] Shouxian Wei, Xiaolin Zheng, Deren Chen, and Chaochao Chen. 2016. A hybrid approach for movie recommendation via tags and ratings. *Electronic Commerce Research and Applications* 18 (2016), 83–94.
- [32] Chenyi Zhang, Ke Wang, Ee-Peng Lim, Qinneng Xu, Jianling Sun, and Hongkun Yu. 2015. Are Features Equally Representative? A Feature-Centric Recommendation. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, Blai Bonet and Sven Koenig (Eds.). AAAI Press, 389–395. <http://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9287>
- [33] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. 2019. Deep Learning Based Recommender System: A Survey and New Perspectives. *ACM Comput. Surv.* 52, 1 (2019), 5:1–5:38. <https://dl.acm.org/citation.cfm?id=3285029>
- [34] Zhou Zhao, Qifan Yang, Hanqing Lu, Tim Wenginger, Deng Cai, Xiaofei He, and Yueting Zhuang. 2017. Social-Aware Movie Recommendation via Multimodal Network Learning. *IEEE Transactions on Multimedia* (2017).