Improving On-line Genre-based Viewer Profiling

Bruno Veloso
University of Vigo
Vigo, Spain
INESC TEC
Porto, Portugal
bruno.miguel.veloso@gmail.com

Benedita Malheiro
ISEP/Polytechnic of Porto
Porto, Portugal
INESC TEC
Porto, Portugal
mbm@isep.ipp.pt

Juan Carlos Burguillo
University of Vigo
Vigo, Spain
J.C.Burguillo@det.uvigo.es

Jeremy Foss
Birmingham City University
Birmingham, United Kingdom
Jeremy.Foss@bcu.ac.uk

ABSTRACT

Typically, recommendation algorithms are unable to make recommendations for new users due to the inherent lack of information, i.e., the cold start problem. To overcome this problem, this work addresses builds new viewer profiles by combining general and personal feature-based profiles using both the frequency and the rating of each feature. For each newly arrived viewer, we create a dynamic profile by combining the corresponding demographic stereotypical profile with the individual profile and, then, as the number of the viewer-generated events increases, we gradually fade the general component and strengthen the individual component. Specifically, we combine the genre frequency & rating of the viewer personal and demographic stereotype profiles. This novel viewer profiling algorithm was evaluated with the MovieLens 100k and 1M data sets, using content-based and collaborative stream mining recommendation techniques. When compared with the standard average user stereotype, the results with the demographic stereotypes show a significant improvement in terms of classification accuracy, identical prediction accuracy and an increase in run time.

CCS CONCEPTS

• **Information systems → Data mining**;

KEYWORDS

User/Viewer Profiling; Personalisation; Recommendation

ACM Reference format:
Bruno Veloso, Benedita Malheiro, Juan Carlos Burguillo, and Jeremy Foss. 2017. Improving On-line Genre-based Viewer Profiling. In *Proceedings of the In-Programme Personalisation for Broadcast (IPP4B) Workshop, ACM TVX 2017 Conference, Hilversum, The Netherlands, June 2017 (TVX 2017),* 6 pages.

Proceedings of the In-Programme Personalisation for Broadcast (IPP4B) Workshop, ACM TVX2017, Hilversum, The Netherlands, June 2017.

1 INTRODUCTION

The goal of this research is to improve the quality of the recommendations provided to new user by exploring any user related features present in the data set. In this work, we experiment with the available user demographic data – gender, age and occupation – to overcome the lack of information on the new user and, thus, improve the accuracy of the recommendations provided to new users.

The main challenge of entity modelling/profiling is to represent as distinctively as possible the different entities based on the available data, which may include explicit (user provided) and implicit (behaviour based) information. In particular, profiles can be based on: (*i*) the frequency of the items – item frequency (IF); (*ii*) the rating of the items – item rating (IR); (*iii*) the frequency of the features – feature frequency (FF); and (*iv*) the rating of the features – feature rating (FR). When compared with item-based profiles, feature-based profiles typically scale-up well, since they produce considerably lower dimension models, and are less sensitive to cold start problems, but tend to generate less accurate recommendations. Standard profiling algorithms are, in the case of content-based filtering, frequency-based (IF or FF), whereas, in the case of collaborative filtering, are classification-based (IR or FR). These algorithms, which are exclusively based on the user generated events (cliques, reviews, ratings, tags, etc.), are unable to provide recommendations involving new entities – users or items – since, by default, they are event-less. This work addresses the so-called new user cold start problem with the temporary help of demographic stereotypes and adopts a novel feature frequency & rating (FFR) based profiling.

This paper is organized in five sections. Section 2 introduces the concept of content-based filtering and collaborative filtering and presents the different profiling techniques, including stereotypes. Section 3 describes the proposed profiling system. Section 4 details the datasets, evaluation protocol, evaluation metrics, tests and discusses the results obtained. Finally, Section 5 draws the conclusions and identifies future developments.

2 RELATED WORK

Recommendation systems rely on the user and item profiles to generate personalised recommendations. Internally, they may implement content-based, collaborative or hybrid filters. Content-based recommendation algorithms recommend items according to the similarity between the user and item profiles. This approach is applicable when the available data holds the collection of items viewed per user, the items rated per user or the items rated and tagged per user. While, as a whole, Content-based Filtering (CbF) suffers from the over-specialization problem, it minimises the item cold-start problem present in collaborative filtering. Collaborative recommendation algorithms are applicable whenever the data set includes the ratings of the items viewed by the viewer or the ratings of the items tagged by the viewer. In particular, k-NN collaborative filters (CF), first, determine the correlation between users or between items and, then, make predictions based on the set of nearest neighbours. For example, when processing just ratings, the correlation between two users can be based on the number of co-rated items – user-based CF – or on the number of users which co-rated the same items – item-based CF.

Features, cliques, reviews, tags and ratings have been used to build individual profiles as well as group profiles. Profiles are typically represented by Vector Support Models (VSM) and built from demographic, textual (reviews, keywords or tags), item or feature related (frequency and rating) information. Whenever there is insufficient data to characterize entities and the number and/or dimension of the profiles increases significantly, group profiling is adopted. Group profiles aggregate under a single profile multiple entities and can be created using clustering or stereotyping. Whereas in the first case, profiles are based on the properties of the objects within the cluster, on the latter case, profiles are based on the properties of objects of a given pre-defined category. In the case of recommendation systems, both group profiling techniques are used to represent new users and new items, *i.e.*, addresses the cold start problem. Typically, stereotypes are feature-based models [1, 14, 15, 18].

Diverse demographic features has been explored by researchers: age and gender [2, 8, 16, 17]; age, gender and nationality [22]; age, gender, ethnicity and political orientation [21]; age, gender and occupation [4, 5, 9, 12]; age (children & adults), gender, level of education and place of living [10]; and age, gender, area code, level of education and employment status [11]. Specifically, this information has been used to create group profiles [16], demographic stereotypes [2, 4] and individual profiles [5, 8–12, 17, 22]. In terms of profiling approaches, demographic and user event data have been employed to create FF profiles [2, 5, 12, 17], FR profiles [8–11, 13, 16] and both IR and FF profiles [4, 22]. In particular, [4, 22] implement item-based collaborative algorithms which determine: (*i*) the relative importance of items and the relevance of the features based, respectively, on IR and FF [4]; and (*ii*) the nearest neighbours by determining the demographic FF correlation between the active and the remaining users and the items predictions using IR [22]. Furthermore, the demographic profiles have been applied in the pre-recommendation stage, e.g., to choose the nearest neighbours in CF [4, 5, 9, 11, 12, 22], and in the recommendation stage, e.g., to determine the similarity between users and items in CbF [10, 11, 13].

Although our proposal uses the same demographic information as [4, 5, 9, 12], we propose a novel feature-based – feature frequency & rating (FFR) – profiling method by building and combining demographic stereotypes with the individual profiles to minimise the new user cold-start problem. Additionally, we also adapt existing data stream techniques to work with our novel profiling technique.

3 ENTITY PROFILING

The profile of an entity, which is represented in Figure 3, includes an individual component based on the past entity-related events and a stereotypical component used to minimise the cold start problem while the number of entity-related events is reduced.

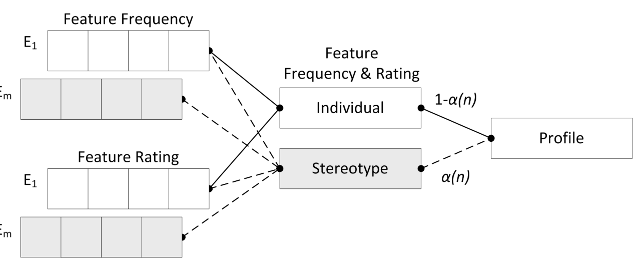

Figure 1: Entity profiling

The combining function αn is defined according to Equation 1 where n represents the number of events related with the entity so far and N the minimum number of events required for a fully individual profile.

$$\alpha n \begin{cases} nN & : n < N \\ 0 & : n \geq N \end{cases} \tag{1}$$

The profile of an entity e, user or item, is given by Equation 2 where n the number of events related with the user or item so far, $\hat{I}_e$ the individual profile, $\hat{S}_e$ the stereotype and $\hat{P}_e$ the current entity profile.

$$\hat{P}_e \ \hat{I}_e \ \alpha n \hat{S}_e - \hat{I}_e \tag{2}$$

Equation 3 presents the incremental update process applied on the viewer profile and stereotypes, where $\hat{P}_e$ is the current entity profile, f is the current feature (genre) and $r_{u,i}$ is the rating defined by the user u for the item i.

$$\hat{P}_e f \ \hat{P}_e f \times \left(\frac{n-1}{n}\right)\left(\frac{r_{u,i}}{n}\right) \tag{3}$$

In our case, the $\hat{P}_e$, $\hat{I}_e$ $\hat{S}_e$ are feature vectors of size twenty, corresponding to the twenty IMDb[1] genres included on the

[1] http://www.imdb.com/genre/

MovieLens dataset. These feature vectors represent the feature frequency & rating (FFR). We experimented with different demographic user stereotypes: Age, Gender, Occupation, Age & Gender, Age & Occupation, Gender & Occupation as well as Age & Gender & Occupation. These stereotypes where built using the genre frequency & rating. We created six different age stereotypes, contemplating six age segments (0-14, 15-24, 25-34, 35-44, 45-54 and 55+), two gender stereotypes (male and female) and 18 different occupation stereotypes, corresponding to the 18 occupational categories present in the data set.

Algorithm 1, which illustrates the designed on-line updating process, requires as inputs the top N recommendations, the minimum number of events for a fully individual user N_u and a fully individual item profile N_i (line 1). The algorithm, for each new data stream event $< u, i, r_{u,i} >$ (line 3), where u represents the user, i the item and $r_{u,i}$ the rating given to the item by the user, calculates the corresponding predicted rating $\hat{r}_{u,i}$ (CF) or similarity (CbF) (line 4), depending on the selected recommendation filter. If the event regards a new user or item, it additionally creates the new entity profile (line 5-6). Then, it evaluates the quality of the recommendation by calculating, in the case of CF, the incremental root mean square error (RMSE) (line 7) and, for both CF and CbF, the incremental TRecall@N and Recall@N (lines 8-15), using a sorted list of 1001 movies composed of 1000 randomly selected unwatched movies plus the newly rated movie (lines 8-15). Line 11 prints the resulting top N recommendations for user u. Finally, it updates the stereotypical and individual profile components followed by the user and item profiles (lines 16-17).

Algorithm 1 On-line Profiling and Recommendation

1: **inputs**: N, N_u, N_i
2: **outputs**: $Recall@N$, $TRecall@N$, $rmse$
3: **for** $r_{u,i} \leftarrow Stream\ Data <u,i,r_{u,i}>$ **do**
4: $\quad \hat{r}_{u,i} \leftarrow CbF\ or\ CF\ prediction$
5: $\quad$ **if** $u \notin Users$ **then** $Create\ \hat{P}_u$
6: $\quad$ **if** $i \notin Items$ **then** $Create\ \hat{P}_i$
7: $\quad$ **if** CF **then** $rmse\ calcRMSE \hat{r}_{u,i}, r_{u,i}$
8: $\quad L_{u,1000} \leftarrow Select\ 1000\ unwatched\ items$
9: $\quad L_{u,1001} \leftarrow \hat{r}_{u,i}$
10: $\quad SL_{u,1001} \leftarrow Sort\ L_{u,1001}$
11: $\quad$ print($SL_{u,1001}$, N)
12: $\quad$ **if** $indexOf \hat{r}_{u,i}, SL_{u,1001} \leq N$ **then**
13: $\quad\quad Recall@N\ updateRecall@N \hat{r}_{u,i}, N$
14: $\quad\quad TRecall@N\ updateTargetRecall \hat{r}_{u,i}, N$
15: $\quad Update\ \hat{S}_u, \hat{S}_i, \hat{I}_u, \hat{I}_i$
16: $\quad Update\ \hat{P}_u, \hat{P}_i$

For the recommendation, we adopted the content-based filter proposed by Veloso *et al.* [19] and the collaborative filter described by Melville *et al.* [7]. The content-based filter adopts the Collinearity and Proximity Similarity metric (CPS) to compute the similarity between users and items. Equation 4 presents the CPS formula, where CS is the cosine similarity, CDD is the Chebyshev distance dissimilarity and β is the combining parameter.

$$\begin{aligned} CPS\ & \beta \times CS\ 1 - \beta \times CDD \\ & \beta \frac{{}^{n}_{j1} \hat{A}_j \hat{B}_j}{\sqrt{{}^{n}_{j1} \hat{A}_j^2}\sqrt{{}^{n}_{j1} \hat{B}_j^2}} \\ & 1 - \beta \times 1 - MAX_j |\hat{A}_j - \hat{B}_j| \end{aligned} \tag{4}$$

The collaborative filter implements a k-NN approach. The algorithm, which uses the Pearson correlation coefficient to compute the similarity between users, selects, for each active user, a subset of users based on their similarity to make predictions regarding the items rated by the neighbours and not yet rated by the active user [7].

4 EXPERIMENTS AND RESULTS

The following subsections describe the performed experiments, including the data sets, the evaluation metrics and protocol, the tests and the results. The experiments were performed with an Intel i7-2600 3.4 GHz Central Processing Unit (CPU) and 16 GB DDR3 Random Access Memory (RAM) platform.

4.1 Data Sets

Our approach was evaluated with two data sets that has demographic data: (i) MovieLens 100k (ML100k)[2]; and (ii) MovieLens 1M[3]. ML100k data set has a data sparsity of 93.7 % and contains information about 943 users and 1682 movies, including 100000 user ratings together with timestamps. These user ratings were collected between the 19th September 1997 and the 22nd April 1998.

ML1M data set has higher data sparsity 95.5 % and size. It contains information about 6040 users and 3952 movies, including 1000000 user ratings together with timestamps. These user ratings were collected between the 26th April 2000 and the 28th February 2003. In both data sets each user has at least 20 movie ratings.

Figure 2 shows the temporal distribution of the new user events in both data sets. The new users are distributed throughout the time span of the data set, excluding for the ML1M dataset the last 10 % of the time span where there are almost no new user events.

The data was ordered temporally and divided in two parts, the first slice with 20 % of the timespan for the train phase and 80 % for the data streaming phase. Our evaluation will be on the data stream timespan to assess the impact of the new users with demographic stereotypes. For the ML100k we have 77.2 % of new users (728 users in 943) and for the ML1M we have 76.3 % of new users (4611 users in 6040).

4.2 Evaluation Metrics

We calculate for each new viewer rating event the Recall@N proposed by Cremonesi et al. [3] and the Target Recall@N

[2]http://files.grouplens.org/datasets/movielens/ml-100k.zip
[3]http://files.grouplens.org/datasets/movielens/ml-1m.zip

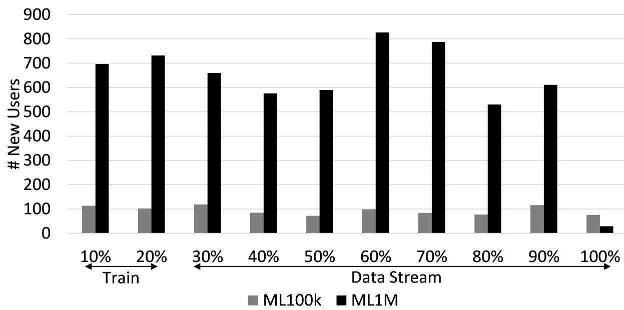

Figure 2: Distribution of new user events in the data sets

presented by Veloso et al. [20]. In the first case, we predict the ratings of all items unseen by the viewer, including the newly rated item, select 1000 unrated items plus the newly rated item and sort them in descending order. If the newly rated item belongs to the list of the top N viewer predicted items, we count a hit. In the second case, we use all rated items instead of just the top-rated items. *i.e.*, the Target Recall@N (TRecall@N) presented by Veloso et al. (2017) [20], evaluates the accuracy of the predictions using all viewer ratings. The TRecall@N metric, which was proposed for the evaluation of collaborative filters, verifies, for each new user rating event, if the predicted rating lays within the short list of N items centred around the actual user rating (target). However, this metric cannot be directly applied to content-based filters since they determine similarities rather than predictions. To circumvent this problem, we compute the similarity value between the target item and the user. First, we multiply each feature in the profile of the target item by the target rating and, then, calculate the similarity between this temporary target item profile with the viewer profile, using CPS. Finally, using this target similarity value, we determine the target recall. This adaptation is a contribution in terms of evaluation metrics since it allows the usage of the same classification metrics for CbF and CF on-line stream mining. In particular, we calculate the Recall@10 (R@10) and the TRecall@10 (TR@10).

Regarding the rating feedback, we determine the global prediction RMSE, which is calculated incrementally after each new viewer rating event.

4.3 Evaluation Protocol

The evaluation protocol defines the data ordering, partitions and distribution. First, the data was ordered temporally and, then, partitioned. The initial off-line model uses the "Batch Train" (20 % of the data set for training). The on-line model uses the "Stream Data", which corresponds to the remaining 80 % of the data set. Each one of these ratings triggers the generation and immediate evaluation of the predictions. In particular, when a viewer rates a movie, the algorithm uses the new rating to update the predictions for that user. Finally, the algorithm updates the viewer, movie, and stereotype profiles. The adopted evaluation method was inspired by the prequential evaluation proposed by Gama *et al.* (2009) [6]. The predictions are evaluated using global RMSE, Recall@10, TRecall@10 metrics. Additionally, we measure the average on-line update time per event. In off-line mode, the average global RMSE and Recall@10 are computed at the end, while in on-line mode, they are calculated whenever a new event occurs.

4.4 Tests

The experiments were conducted both with ML100k and ML1M data sets and involved eight viewer stereotypes: the global average (base algorithm), the Age, the Gender, the Occupation, the Age & Gender, the Age & Occupation, the Gender & Occupation and the Age & Gender & Occupation viewer stereotypes. For each case, the optimal average number of viewer events N to display an exclusively personal viewer profile was determined. Initially, we varied N from 10 till 200, using increments of 10 and, next, explored with unit increments the most promising ranges, namely from 1 till 60, using increments of 1. The selection of the best viewer stereotype was based on the highest TRecall@10 value.

Table 1 and Figure 3 presents the content-based filtering results with the tested stereotypes. In the case of ML100k, the Gen stereotype shows an improvement of 28.6 % in TRecall@10 and 81.1 % in Recall@10 with the same run time. The optimal average number of viewer events N for a fully personal profile was 7. In the case of ML1M, the Age-Gen-Occ stereotype shows an improvement of 21.1 % in TRecall@10, 25.0 % in Recall@10 and a run time increase of 37.5 %. The optimal average number of viewer events N for a fully personal profile was 8.

Table 1: CbF Results with Different Viewer Stereotypes

	Stereotypes	N	TR@10	R@10	Δt (s)
ML100k	Average (base)	20	0.014	0.011	0.7
	Age	5	0.017	0.020	0.6
	Gen	**7**	**0.018**	**0.020**	**0.7**
	Occ	4	0.017	0.021	0.6
	Age-Gen	2	0.017	0.020	0.8
	Age-Occ	2	0.017	0.020	0.7
	Gen-Occ	4	0.016	0.019	0.7
	Age-Gen-Occ	7	0.017	0.019	0.6
ML1M	Average (base)	20	0.019	0.024	0.8
	Age	6	0.022	0.029	0.9
	Gen	8	0.022	0.030	1.0
	Occ	5	0.023	0.028	1.0
	Age-Gen	8	0.023	0.029	1.0
	Age-Occ	5	0.022	0.029	0.9
	Gen-Occ	2	0.022	0.029	1.1
	Age-Gen-Occ	**8**	**0.023**	**0.030**	**1.1**

Table 2 and Figure 4 presents the collaborative filtering results with the different demographic viewer profiles. In the case of ML100k, the best stereotype was the Gender with an improvement of 34.7 % in TRecall@10, 8.7 % in Recall@10, a neglegible RMSE increase of 0.5 % and a run time increase of 50.0 %. The optimal average number of viewer events N for an exclusively personal profile was 46. In the case of ML1M, the best demographic stereotype was Age & Occupation with an improvement of 29.7 % in TRecall@10, 5.9 % in Recall@10,

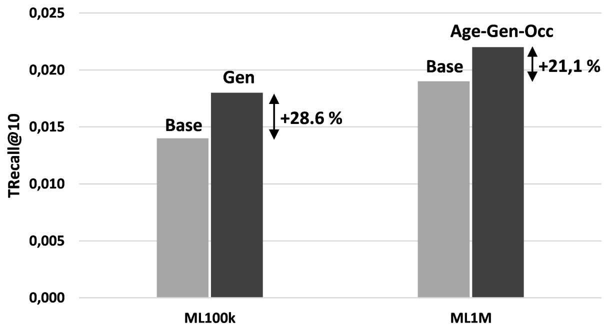

Figure 3: CbF - Results

Table 2: CF Results with Different Viewer Stereotypes

	Stereotypes	N	TR@10	R@10	RMSE	Δt (s)
ML100k	Average (base)	20	0.049	0.023	0.218	0.4
	Age	40	0.060	0.026	0.218	0.6
	Gen	**46**	**0.066**	**0.025**	**0.219**	**0.6**
	Occ	110	0.058	0.029	0.223	0.6
	Age-Gen	35	0.065	0.025	0.219	0.6
	Age-Occ	100	0.060	0.028	0.224	0.6
	Gen-Occ	58	0.063	0.024	0.222	0.6
	Age-Gen-Occ	56	0.066	0.023	0.220	0.6
ML1M	Average (base)	20	0.064	0.034	0.223	1.2
	Age	35	0.073	0.030	0.223	1.2
	Gen	60	0.071	0.029	0.221	1.3
	Occ	35	0.079	0.033	0.223	1.4
	Age-Gen	55	0.075	0.032	0.223	1.3
	Age-Occ	**53**	**0.083**	**0.036**	**0.223**	**1.5**
	Gen-Occ	41	0.081	0.034	0.222	1.3
	Age-Gen-Occ	36	0.076	0.033	0.222	1.3

equal RMSE and a run time increase of 25.0 %. The optimal average number of viewer events N for a fully personal profile was 53.

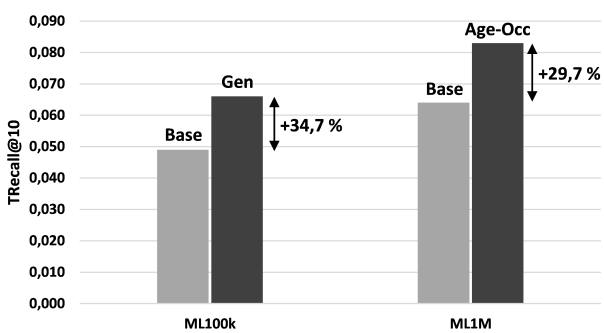

Figure 4: CF - Results

These results clearly show the benefits of our profiling approach: the classification accuracy improves and the prediction accuracy remains unchanged. Moreover, the generated feature-based profiles scale-up well due to their low dimensionality. The execution time growth is proportional to the data set size in the case of CbF and inversely proportional in the case of CF. Although this behaviour requires further research, it suggests that our profiling approach is well suited for on-line data stream mining with CF. In both datasets, the number of viewer events needed to completely fade the demographic stereotypical component is smaller for the CbF than for CF. This means that CF requires more user events to build a fully individual user model than CbF.

5 CONCLUSION

This paper describes a demographic feature-based user profiling algorithm for on-line content-based and memory-based collaborative filtering. The profiles have a stereotypical and an individual component which can be represented using feature frequency & rating.

The proposed algorithm, which uses demographic stereotypes at the early life of the entity profile, helps to reduce the impact of the cold start problem. The results show, in the case of content-based filtering, an increase in the quality of the recommendations with longer run time, while, in the case of collaborative filtering, improved quality of recommendations, identical prediction errors and larger execution time. Specifically, in the case of ML1M, the content-based algorithm took 37.5 % longer and the TR@10 improved 21.1 %, while, in the case of collaborative algorithm, the execution time increased 25.0 % and TR@10 improved 29.7 %.

Concerning future work, we are planning to explore demographic group profiling as well as post-recommendation filtering based on demographic data.

6 ACKNOWLEDGMENTS

This work was partially financed by the European Regional Development Fund (ERDF) through the Operational Programme for Competitiveness and Internationalisation (COMPETE Programme), within project «FCOMP-01-0202-FEDER-023151» and project «POCI-01-0145-FEDER-006961», and by national funds through the Fundação para a Ciência e Tecnologia (FCT) - Portuguese Foundation for Science and Technology - as part of project UID/EEA/50014/2013.

REFERENCES

[1] Liliana Ardissono, Cristina Gena, Pietro Torasso, Fabio Bellifemine, Angelo Difino, and Barbara Negro. 2004. User modeling and recommendation techniques for personalized electronic program guides. In *Personalized Digital Television - Targeting Programs to Individual Viewers, volume 6 of Human-Computer Interaction Series, chapter 1*. Kluwer Academic Publishers, 3–26.

[2] J Ávila, X Riofrío, K Palacio-Baus, M Espinoza-Mejía, and V Saquicela. 2015. Semantic Recommender Systems for Digital TV: From Demographic Stereotyping to Personalized Recommendations. In *Computer Aided System Engineering (APCASE), 2015 Asia-Pacific Conference on*. IEEE, 392–396.

[3] Paolo Cremonesi, Yehuda Koren, and Roberto Turrin. 2010. Performance of recommender algorithms on top-n recommendation tasks. In *Proceedings of the fourth ACM conference on Recommender systems*. ACM, 39–46.

[4] Maria Dima, Dimitrios Vogiatzis, George Paliouras, and Panayiotis Stamatopoulos. 2010. Expert based prediction of user preferences. In *Semantic Media Adaptation and Personalization (SMAP), 2010 5th International Workshop on*. IEEE, 112–117.

[5] Simon Fong, Yvonne Ho, and Yang Hang. 2008. Using genetic algorithm for hybrid modes of collaborative filtering in online recommenders. In *Hybrid Intelligent Systems, 2008. HIS'08. Eighth International Conference on*. IEEE, 174–179.

[6] João Gama, Raquel Sebastião, and Pedro Pereira Rodrigues. 2009. Issues in evaluation of stream learning algorithms. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 329–338.

[7] Prem Melville and Vikas Sindhwani. 2010. Recommender systems. In *Encyclopedia of Machine Learning*. Springer, New York, NY, USA, 829–838.

[8] Mona Nasery, Matthias Braunhofer, and Francesco Ricci. 2016. Recommendations with optimal combination of feature-based and item-based preferences. In *Proceedings of the 2016 Conference on User Modeling Adaptation and Personalization*. ACM, 269–273.

[9] An-Te Nguyen, Nathalie Denos, and Catherine Berrut. 2007a. Improving new user recommendations with rule-based induction on cold user data. In *Proceedings of the 2007 ACM conference on Recommender systems*. ACM, 121–128.

[10] Ngoc Thanh Nguyen, Maciej Rakowski, Michal Rusin, Janusz Sobecki, and Lakhmi C Jain. 2007b. Hybrid filtering methods applied in web-based movie recommendation system. In *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*. Springer, 206–213.

[11] Michael J Pazzani. 1999. A framework for collaborative, content-based and demographic filtering. *Artificial intelligence review* 13, 5-6 (1999), 393–408.

[12] Olivier Potonniée. 2004. A decentralized privacy-enabling TV personalization framework. In *European Conference on Interactive Television: Enhancing the Experience (euroITV)*.

[13] Lingyun Qiu and Izak Benbasat. 2010. A study of demographic embodiments of product recommendation agents in electronic commerce. *International Journal of Human-Computer Studies* 68, 10 (2010), 669–688.

[14] Elaine Rich. 1979. User modeling via stereotypes. *Cognitive science* 3, 4 (1979), 329–354.

[15] Guy Shani, Amnon Meisles, Yan Gleyzer, Lior Rokach, and David Ben-Shimon. 2007. A stereotypes-based hybrid recommender system for media items. In *Workshop on Intelligent Techniques for Web Personalization, Vancouver, Canada*. 76–83.

[16] Sven Ewan Shepstone, Zheng-Hua Tan, and Søren Holdt Jensen. 2013. Demographic recommendation by means of group profile elicitation using speaker age and gender recognition.. In *INTERSPEECH*. 2827–2831.

[17] Janusz Sobecki, Emilia Babiak, and M Słanina. 2006. Application of hybrid recommendation in web-based cooking assistant. In *Knowledge-Based Intelligent Information and Engineering Systems*. Springer, 797–804.

[18] Mikael Sollenborn and Peter Funk. 2002. *Category-Based Filtering and User Stereotype Cases to Reduce the Latency Problem in Recommender Systems*. Springer Berlin Heidelberg, Berlin, Heidelberg, 395–405. DOI:http://dx.doi.org/10.1007/3-540-46119-1_29

[19] Bruno Veloso, Benedita Malheiro, and Juan Carlos Burguillo. 2015. A Multi-Agent Brokerage Platform for Media Content Recommendation. *International Journal of Applied Mathematics and Computer Science* 25, 3 (2015), 513–527.

[20] Bruno Veloso, Benedita Malheiro, Juan C. Burguillo, and Jeremy Foss. 2017. Personalised Fading for Stream Data. In *SAC 2017: Symposium on Applied Computing Proceedings, 32nd ACM Symposium on Applied Computing (SAC 2017), Data Streams Track*. ACM, New York, 1–3. To appear.

[21] Udi Weinsberg, Smriti Bhagat, Stratis Ioannidis, and Nina Taft. 2012. BlurMe: Inferring and obfuscating user gender based on ratings. In *Proceedings of the sixth ACM conference on Recommender systems*. ACM, 195–202.

[22] Billy Yapriady and Alexandra L Uitdenbogerd. 2005. Combining demographic data with collaborative filtering for automatic music recommendation. In *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*. Springer, 201–207.