On the Use of Cross-Domain User Preferences and Personality Traits in Collaborative Filtering

Ignacio Fernández-Tobías(✉) and Iván Cantador

Universidad Autónoma de Madrid, 28049, Madrid, Spain
{ignacio.fernandezt,ivan.cantador}@uam.es

Abstract. We present a study comparing collaborative filtering methods enhanced with user personality traits and cross-domain ratings in multiple domains on a relatively large dataset. We show that incorporating additional ratings from source domains allows improving the accuracy of recommendations in a different target domain, and that in certain cases, it is better to enrich user models with both cross-domain ratings and personality trait information.

Keywords: Collaborative filtering · Personality · Cross-domain recommendation

1 Introduction

Most recommendation services exploit user preferences obtained explicitly (e.g., by means of ratings) or implicitly (e.g., by mining click-through and log data). Effective hybrid recommendation approaches have been proposed that also exploit auxiliary data, such as user demographics, item metadata, and contextual signals. Recently, new sources of side information have been explored to enrich user models for collaborative filtering (CF). In particular, it has been shown that people with similar personality traits are likely to have similar preferences [3, 9], and that correlations between user preferences and personality traits allow improving personalized recommendations [8, 12]. Moreover, cross-domain recommendation methods [4] have been shown to be effective in target domains, by exploiting user preferences in other source domains [1, 2, 10]. Previous studies have investigated these sources of auxiliary information, focusing on particular approaches and domains, and in general using relatively small datasets. In this paper, we evaluate various CF methods enhanced with user personality traits and cross-domain ratings. Our empirical results on 22,289 Facebook user profiles with preferences for items in several domains –movies, TV shows, music and books– show that incorporating additional ratings from other domains improves recommendation accuracy, and that in certain cases, it is better to enrich user models with both cross-domain rating and personality trait information.

2 Related Work

2.1 User Personality in Personalized Services

Among the different models proposed to represent human personality, the Five Factor model (FFM) is considered one of the most comprehensive, and has been the mostly used to build user personality profiles [8]. This model establishes five dimensions

F. Ricci et al. (Eds.): UMAP 2015, LNCS 9146, pp. 343–349, 2015.
DOI: 10.1007/978-3-319-20267-9_29

(traits) to describe personality: *openness*, *conscientiousness*, *extraversion*, *agreeableness* and *neuroticism*. Recent research has shown that correlations between user preferences and personality factors exist in certain domains [3, 9], and that these correlations can be used to enhance personalized recommendations [7]. For instance, Hu and Pu [8] presented a method in which user similarities are computed as the Pearson's coefficient of their FF scores, and combined this approach with rating-based similarities to improve CF. Tkalčič et al. [12] evaluated three user similarity metrics for heuristic-based CF, and showed that approaches using FF data perform statistically equivalent or better than rating-based approaches, especially in cold-start situations.

In this paper, we also integrate personality information into the user-based nearest-neighbors algorithm, but we explore additional similarity functions besides Pearson's correlation, and evaluate the quality of the recommendations in several domains.

2.2 Cross-Domain Recommendation

Cross-domain recommender systems aim to generate personalized recommendations in a target domain by exploiting knowledge from source domains [4]. This problem has been addressed by means of user preference aggregation in user modeling [1, 2, 10], as a potential solution to the cold-start and sparsity problems in recommender systems [11], and as a practical application of knowledge transfer in machine learning [6]. We distinguish between two main types of approaches: those that *aggregate* knowledge from various source domains, for example, user preferences [1], user similarities [10], and rating estimations [2], and those that *link* or *transfer* knowledge between domains to support recommendations in a target domain [11].

In this paper, we analyze how *aggregating* ratings from different domains can help improving recommendations in a target domain. Furthermore, we empirically compare cross-domain user preferences and personality traits as valuable sources of auxiliary information to enhance heuristic-based CF methods. Approaches based on knowledge transfer are postponed for future investigation.

3 Integrating Personality in Collaborative Filtering

As done in [8], we integrate personality information into heuristic-based CF methods via the user similarity function. Specifically, we focus on the user-based k nearest-neighbors method for unary/binary user feedback (*likes*, *thumbs up/down*), in which ratings are predicted as:

$$\hat{r}(u,i) = \sum_{v \in N(u)} sim_{pref}(u,v) \cdot \mathbb{I}\big(i \in I(v)\big) \tag{1}$$

where $N(u)$ is the user's neighborhood, $I(v)$ is the set of items rated by user v, and the function $\mathbb{I}(p) = 1$ if p is true, and 0 otherwise. A popular user similarity function in this setting is the Jaccard's coefficient:

$$sim_{pref}(u,v) = \frac{|I(u) \cap I(v)|}{|I(u) \cup I(v)|} \tag{2}$$

This metric is considered as baseline in our experiments. To complement it with user personality information we use a linear combination, analogously to [8]:

$$sim(u,v) = \lambda \cdot sim_{pref}(u,v) + (1-\lambda) \cdot sim_{pers}(u,v) \quad (3)$$

where $\lambda \in [0,1]$ controls the influence of user preferences and personality on the recommendation process. For λ values close to 1, user preferences are more relevant, while for λ values close to 0, personality profiles get higher relevance.

We study several formulations of $sim_{pers}(u,v)$ that yield the compared personality-based CF methods, namely cosine similarity (COS), Pearson's correlation (PEA), Spearman's correlation (SPE), and Kendall's correlation (KEN). When combined with the preference-based similarity as in equation (2), we call the methods as COS-λ, PEA-λ, SPE-λ, and KEN-λ. When $\lambda = 0$, we use COS-pers (respectively PEA-pers, etc.) to name the methods, since only personality information is used in the computation of the user similarity. We refer to [5] for more details on the implementation of the above similarities.

4 Integrating Cross-Domain Ratings in Collaborative Filtering

Cross-domain recommender systems aim to exploit knowledge from source domains $\mathcal{D}_S$ to perform recommendations in a target domain $\mathcal{D}_T$. Without loss of generality, we can consider two domains $\mathcal{D}_S$ and $\mathcal{D}_T$ –the definitions are extensible to more source domains. Let U_S and U_T be their sets of users, and let I_S and I_T be their sets of items. The users of a domain are those who expressed preferences (e.g., ratings, tags) for the domain items.

In our study, we are interested in comparing the effects of cross-domain ratings and personality traits as auxiliary user information in the CF framework. Hence, we focus on the data itself and use the same recommendation algorithm as in Section 3, user-based k nearest-neighbors, in order to provide a fair comparison. Assuming $I_S \cap I_T = \emptyset$, we distinguish between two different scenarios of user overlap:

- *User overlap*. There are some common users who have preferences for items in both domains, i.e., $U_S \cap U_T \neq \emptyset$. This is the case, for instance, where some users rated both movies and books. Recommendations of items in $\mathcal{D}_T$ are generated exploiting user similarities based on preferences for items in $\mathcal{D}_S$ and $\mathcal{D}_T$:

$$sim_{pref}(u,v) = \frac{|I_S(u) \cap I_S(v)| + |I_T(u) \cap I_T(v)|}{|I_S(u) \cup I_S(v)| + |I_T(u) \cup I_T(v)|} \quad (4)$$

- *No overlap*. There is no overlap between users and items in the domains, i.e., $U_S \cap U_T = \emptyset$. Recommendations of items in $\mathcal{D}_T$ are generated by exploiting user similarities based only on user preferences for items in $\mathcal{D}_S$.

$$sim_{pref}(u,v) = \frac{|I_S(u) \cap I_S(v)|}{|I_S(u) \cup I_S(v)|} \quad (5)$$

In both cases, recommendations are generated for users in the target domain $\mathcal{D}_T$.

5 Experiments

5.1 Dataset

The dataset used in our experiments was obtained by the myPersonality project (`http://mypersonality.org`). Due to the size and complexity of the database, in this paper we restrict our study to a subset of its items. Specifically, we selected all likes (ratings) associated to items belonging to movie genres, music genres, book genres, and TV genres. Thus, for instance, selected items belonging to the movie genre category are `comedy`, `action`, `star wars`, and `james bond`. Table 1 shows some statistics about the users, items and ratings in the four considered domains.

Table 1. Statistics of the used dataset

domains	users	items	ratings	rating sparsity	user overlap
movies	16,168	268	27,921	99.36%	N/A
music	17,980	1,175	66,079	99.69%	N/A
books	15,251	305	23,882	99.49%	N/A
TV	4,142	111	4,612	98,99%	N/A
movies + music	22,012	1,443	94,000	99.70%	55.13%
movies + books	21,410	573	51,803	99.58%	46.75%
movies + TV	17,671	379	32,533	99.51%	14.93%
music + books	22,029	1,480	89,961	99.72%	50.85%
music + TV	19,201	1,286	70,691	99.71%	15.21%
books + TV	16,766	416	28,494	99.59%	15.67%

5.2 Evaluation Setting

In our experiments we empirically compared the performance of the user-based CF method extended with user cross-domain ratings and personality traits, respectively. We also tested different single-domain baselines in order to analyze the effect of the additional data on the quality of the recommendations.

- **Most popular**. Non-personalized approach that recommends the most liked items.
- **iMF**. Matrix factorization for positive-only feedback. The number of latent factors is set to 10, as we did not observe significant differences with respect to other values in preliminary tests.
- **Item kNN**. Item-based nearest neighbors using Jaccard's similarity. In our experiments, we set $k = \infty$ and considered all similar items rated by the target user.
- **User kNN**. User-based nearest neighbors using Jaccard's similarity. In our experiments, we set $k = 50$. When extended with cross-domain preferences, we used equations (4) and (5) to compute user similarities.
- **Personality-based CF**. User-based nearest neighbors extended with user personality information, using different instances of the similarity in equation (3) with COS, PEA, SPE, and KEN, and $\lambda \in \{0, 0.1, 0.2, \dots, 0.9\}$. Note that $\lambda = 1$ corresponds to standard User kNN. We also combined these methods with cross-domain preferences, inserting similarities in equations (4) and (5) into equation (3).

These methods were evaluated in terms of MAP, F@5, and coverage, measured as the number of users for which recommendations could be generated. All results were averaged using 5-fold cross validation. For both single- and cross-domain recommendations, the test sets were composed of ratings for items in the target domain, and were never used for computing user similarities.

5.3 Results

Table 2 shows the results of the best performing baseline and personality-based methods (in terms of MAP), in both user overlap and non-overlap cross-domain scenarios. Results for single-domain recommendation are reported in rows where source and target domains match. The best values for each scenario are in bold, and the overall best for each target domain are underlined. Significant differences (Wilcoxon signed rank test, $p < 0.05$) are marked as follows: ▲ against single-domain baseline, * against cross-domain User kNN with same source and overlap configurations, and † for user overlap against no overlap, i.e., by rows.

Table 2. Recommendation performance in single- and cross-domain scenarios

Target domain	Source domain	No overlap				User overlap			
		Method	MAP	F@5	Coverage	Method	MAP	F@5	Coverage
Books	Books	Item kNN	0.2822	0.1801	**0.4353**				
		PEA-0.1	0.2679	0.1632	0.4175				
	Movies	User kNN	0.4859▲	0.2387▲	0.3567	User kNN	0.4907▲	0.2271▲	0.3189
		KEN-0.8	**0.4927***	**0.2429***	0.3565	KEN-0.9	**0.4952**	**0.2288**	0.3193
	Music	User kNN	0.2526	0.1669	0.3362	User kNN	0.2674†	0.1684	0.3311
		SPE-0.9	0.2565	0.1674	0.3360	PEA-0.7	0.2736*†	0.1716*†	**0.3348**
	TV	User kNN	0.3680▲	0.2038▲	0.1095	User kNN	0.3830▲†	0.2029▲	0.0980
		COS-0.9	0.3726	0.2084	0.1095	COS-0.9	0.3847	0.2037	0.0979
Movies	Movies	Item kNN	0.5303	0.2716	**0.4994**				
		KEN-0.1	0.5341	0.2655	0.4678				
	Books	User kNN	**0.6940**▲	0.3228▲	0.3853	User kNN	**0.6845**▲	**0.3122**▲	0.3412
		COS-0.2	0.6918	**0.3229**	0.3853	COS-0.9	0.6817	0.3115	0.3412
	Music	User kNN	0.5311	0.2797▲	0.4030	User kNN	0.5502▲†	0.2715	0.3941
		COS-0.1	0.5343	0.2806	0.4030	SPE-0.5	0.5587†	0.2820*†	**0.4024**
	TV	User kNN	0.5530	0.2776	0.1048	User kNN	0.5785▲†	0.2785	0.0933
		SPE-0.9	0.5757	0.2893*	0.1047	PEA-0.8	0.5856†	0.2844*	0.0945
Music	Music	Item kNN	**0.3225**	**0.2335**	**0.5348**				
		KEN-0.9	0.3054	0.2249	0.4293				
	Books	User kNN	0.2207	0.1736	0.3458	User kNN	0.3032†	0.2207†	0.3424
		COS-0.3	0.2277*	0.1780*	0.3458	COS-0.8	0.3036†	0.2207†	0.3424
	Movies	User kNN	0.2504	0.1890	0.3871	User kNN	0.3201†	0.2338†	0.3802
		SPE-0.8	0.2541*	0.1920*	0.3867	COS-0.9	**0.3209**†	**0.2339**†	**0.3803**
	TV	User kNN	0.1880	0.1572	0.0942	User kNN	0.2651†	0.2070†	0.0941
		COS-0.7	0.1971*	0.1646*	0.0942	COS-0.8	0.2661†	0.2082†	0.0941
TV	TV	Most popular	0.4174	**0.2320**	**0.9766**				
		COS-pers	0.3686	0.2106	0.1492				
	Books	User kNN	0.4375	0.2143	0.1201	User kNN	0.4426	0.1961	0.1045
		PEA-0.8	0.4561	0.2176	0.1201	KEN-0.9	0.4591	0.2047	0.1043
	Movies	User kNN	0.4322▲	0.2022	0.1075	User kNN	**0.4645**▲†	0.1998	0.0891
		PEA-0.9	**0.4682***	0.2217*	0.1075	PEA-0.8	0.4618	0.2076	0.0915
	Music	User kNN	0.3882	0.2248	0.1160	User kNN	0.3537	0.1966	0.1123
		COS-0.3	0.4073*	0.2284	0.1160	SPE-0.4	0.3817	**0.2191***	**0.1160**

In the **books** domain, the best personality-based method PEA-0.1 performs worse than the best baseline, Item kNN. Methods exploiting cross-domain preferences from movies and TV shows achieve better performance, although with the latter the coverage drops drastically. We find that this behavior is repeated in the other target domains, and argue that it is due to their low user overlap with the TV shows domain (Table 1). Combining personality traits and cross-domain ratings slightly improves the performance. In the **movies** domain, the improvement achieved by personality is not significant enough, while cross-domain preferences, especially from the books domain, provide much better precision. In this case, however, the user coverage drops roughly 10%. In the **music** domain, the best single-domain baseline is unbeaten, and neither personality nor cross-domain ratings are useful. This seems to indicate that the users' choices on music are only determined by their preferences on this domain. Finally, in **TV shows** domain the same trend as with books and movies is observed: the popularity baseline outperforms the best personality-based method in single-domain recommendation, whereas cross-domain ratings prove to be more valuable. In this case, in contrast, it would be best to stick with the baseline, since it offers a much better coverage and overall tradeoff with accuracy.

In general, even though personality does not outperform a strong single-domain baseline, it allows for further improving cross-domain methods. Also, results in the user overlap cross-domain scenario are just slightly better than in the non-overlap case. This is somewhat unexpected and counterintuitive, and we argue that it may be due to the fact that items represent genres (not items) in our dataset, and users with similar preferences in the source domain are likely to also share similar preferences in the target domain, with possibly the exception of music, as we conjectured before.

On a final note, Item kNN consistently outperformed iMF and User kNN as single-domain baselines. Again, this is likely caused by the fact that in our dataset we are dealing with genres, which can be considered as very popular items with many ratings. Item kNN takes advantage of this by computing more robust item to item similarities.

6 Conclusions and Future Work

We have preliminary compared the suitability of user cross-domain preferences and personality traits as auxiliary information to improve CF. We evaluated a number of personality-aware and cross-domain methods on the top-N recommendation task in various domains. From the achieved empirical results, we conclude that in general it is more valuable to collect user preferences from related domains than personality information to improve the quality of book and movie recommendations. Nonetheless, user personality can be exploited to further enhance cross-domain methods, most notably in the TV domain.

We plan to extend our study using larger datasets. In this paper we tested a large amount of user personality profiles, but focused on preferences for item genres. We want to analyze whether personality and cross-domain ratings can improve recommendation performance when dealing with particular movies, TV shows, music albums, or books.

Acknowledgements. This work was supported by the Spanish Ministry of Economy and Competitiveness (TIN2013-47090-C3). We thank Michal Kosinski and David Stillwell for their attention regarding the dataset.

References

1. Abel, F., Helder, E., Houben, G.-J., Henze, N., Krause, D.: Cross-system User Modeling and Personalization on the Social Web. UMUAI **23**(2–3), 169–209 (2013)
2. Berkovsky, S., Kuflik, T., Ricci, F.: Mediation of User Models for Enhanced Personalization in Recommender Systems. UMUAI **18**(3), 245–286 (2008)
3. Cantador, I., Fernández-Tobías, I., Bellogín, A.: Relating Personality Types with User Preferences in Multiple Entertainment Domains. EMPIRE 2013 (2013)
4. Fernández-Tobías, I., Cantador, I., Kaminskas, M., Ricci, F.: Cross-domain Recommender Systems: A Survey of the State of the Art. CERI **2012**, 187–198 (2012)
5. Fernández-Tobías, I., Cantador, I.: Personality-aware collaborative filtering: an empirical study in multiple domains with facebook data. In: Hepp, M., Hoffner, Y. (eds.) EC-Web 2014. LNBIP, vol. 188, pp. 125–137. Springer, Heidelberg (2014)
6. Gao, S., Luo, H., Chen, D., Li, S., Gallinari, P., Guo, J.: Cross-domain recommendation via cluster-level latent factor model. In: Blockeel, H., Kersting, K., Nijssen, S., Železný, F. (eds.) ECML PKDD 2013, Part II. LNCS, vol. 8189, pp. 161–176. Springer, Heidelberg (2013)
7. Hu, R., Pu, P.: A study on user perception of personality-based recommender systems. In: De Bra, P., Kobsa, A., Chin, D. (eds.) UMAP 2010. LNCS, vol. 6075, pp. 291–302. Springer, Heidelberg (2010)
8. Hu, R., Pu, P.: Enhancing Collaborative Filtering Systems with Personality Information. RecSys **2011**, 197–204 (2011)
9. Rentfrow, P.J., Goldberg, L.R., Zilca, R.: Listening, Watching, and Reading: The Structure and Correlates of Entertainment Preferences. Journal of Personality **79**(2), 223–258 (2011)
10. Shapira, B., Rokach, L., Freilikhman, S.: Facebook Single and Cross Domain Data for Recommendation Systems. UMUAI **23**(2–3), 211–247 (2013)
11. Shi, Y., Larson, M., Hanjalic, A.: Tags as bridges between domains: improving recommendation with tag-induced cross-domain collaborative filtering. In: Konstan, J.A., Conejo, R., Marzo, J.L., Oliver, N. (eds.) UMAP 2011. LNCS, vol. 6787, pp. 305–316. Springer, Heidelberg (2011)
12. Tkalčič, M., Kunaver, M., Košir, A., Tasič, J.F.: Addressing the New User Problem with a Personality Based User Similarity Measure. UM4Motivation 2011 (2011)