Personalized Video Recommendations for Shared Accounts

Shuo Yang
School of Informatics and Computing
Indiana University
shuoyang@indiana.edu

Somdeb Sarkhel
Adobe Research
San Jose
sarkhel@adobe.com

Saayan Mitra
Adobe Research
San Jose
smitra@adobe.com

Viswanathan Swaminathan
Adobe Research
San Jose
vishy@adobe.com

***Abstract*—Account sharing is a significant problem for online recommender systems to generate accurate personalized recommendations. To solve this problem, one not only has to identify whether an account is shared, but also needs to recognize the different users sharing that account. However, to generate relevant, personalized recommendations, the particular user under a shared account has to be correctly identified at the time of delivering the recommendations. In this paper, we address this problem by first identifying users behind each account using a projection based unsupervised method, and then learning a function which can predict a user's preference accurately based on their 'contextual' information. This approach allows us to generate personalized recommendations for each of the users sharing a single account. We empirically show that on real and synthetic data set our approach performs better than other state-of-the-art approaches.**

I. INTRODUCTION

With the assistance of recommendation systems, the Internet protocol television service has greatly improved the user experience by efficiently predicting user preferences from viewing logs. However, the prediction accuracy could degrade significantly when there are multiple individuals sharing the same user account, especially when their interests are markedly different from each other. Persona identification within a shared user account is a challenging problem, but solving this can greatly influence the performance of the video recommendations, advertisement campaigns or any other personalized recommender system.

We aim to detect the existence of multiple personas [1] behind each user account and identify different preferences associated with every single person sharing this account. Our hypothesis is that the individuals sharing the same account have different video preferences and such variation of preferences depends on user-context features which we defined as the user-related context information at each time a video session has been watched and recorded, such as location, time, device, etc. For any use case which is consistent with such assumption, our proposed system would be able to probabilistically identify the persona based on the contextual features. Hence, the recommender could make use of such information and provide more personalized recommendations for individuals who share a common account.

[1]we use persona and person interchangeably as without loss of generality the solution applies to a single person with multiple personas too.

II. RELATED WORK

Video recommendations are often made in a platform (such as IPTV) suited for group viewing or where accounts are shared among multiple individuals. Identifying which individual is consuming the content at a particular time or context is paramount to producing personalized recommendations. Recommending videos without visibility into the accounts or groups is challenging [1], [2]. Wang et al. [3] tackle the problem by introducing virtual users who represent activities from a single account in a sub-period. Each account is split into virtual users and then similar virtual users are merged to identify persona. They assume that different persona in an account have different viewing patterns according to time of the day. In [4], Yang et al. also identify persona based on the temporal viewing pattern of different genre. Similar viewing patterns in a single time interval are attributed to the same user. In [5], the authors employ a similar approach and use logs which indicate times when the television was switched on or off and the behavior is analyzed for repeating daily/weekly patterns. The common technique of identifying users based on time is evaluated by the authors in [6]. The authors in [7] have a similar approach to identify persona. They first group similar channels and then attribute them to individuals in an account. In [8], the authors propose a model to map movies to users in an household in the latent space and suggest a statistical test to test the composites of the household account. In [9], the authors treat the input data which only has binary positive feedback. Instead of identifying which constituent users are active at a particular time, they pick the top-N recommendations which will most satisfy all the users sharing the account and present them with an explanation on why each item was picked.

III. PERSONA IDENTIFICATION SYSTEM

We propose a 'persona' identification system which exploits rich contextual information (e.g., device, location, time) present in user viewing logs. We define 'persona' as distinct, different preferences present in an account. The primary reason for the presence of multiple preference is account sharing. However, it is possible that a single account holder is showing different preferences under different situations.

There are two main components in our system, *Persona Detection* and *Persona Prediction*. The overall architecture

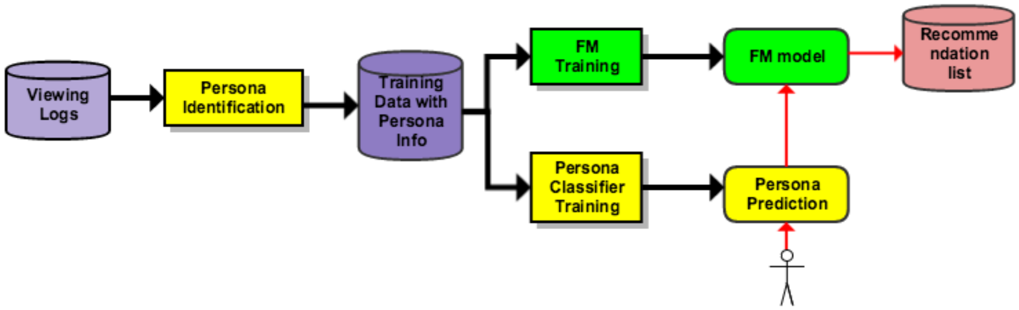

Fig. 1. Persona detection framework.

of our approach is shown in Fig 1, and the aforementioned components are colored in yellow. We used a state-of-the-art recommender system, the Factorization Machines (FM) [10] for generating the recommendations. The goal of the persona detection is to discover whether the viewing log data contains a pattern implying dissimilar preferences within a given account. Whereas the goal of the persona prediction is to analyze whether there is a correlation between a specific persona and its corresponding contextual features (e.g. location, time, device, etc.) and if so, how to accurately predict the persona based on such context information. In the next two subsections we described these two components in greater details.

A. Persona Detection

To detect whether multiple personas exist in an account we use a projection based method.

Definition 1: [*Projection*] A projection ϕ of a given set $x_1, ..., x_m$ in $\mathbb{R}^n$ is a mapping to the set $y_1, ..., y_m$ in $\mathbb{R}^d$, i.e. $y_i = \phi(x_i)$.

Here, we are interested in a ϕ such that $d << n$ and y_i are representative of x_i. The most popular choice for calculating ϕ is through *Principal Component Analysis* (PCA). Other popular choice of projection techniques include *Locality Preserving Projections*, *Isomap*, *Autoencoders*, etc. In this paper, we use PCA as our projection technique, with $d = 1$.

Our approach for persona detection is as follows. First we decide on a suitable feature vector representation of an user rating event. If explicit rating information is unavailable we will use the 'session progress' as the rating by a user for a video as suggested by [11]. Then we project each of these rating events into a lower dimension. This makes it more approachable to compare different rating events belonging to a single account. We observe that if there is a single persona in an account the lower dimensional projections of the rating events will be close to each other, whereas they will be farther apart if multiple personas are present. Hence, we can identify distinct personas by clustering the rating event points in this dimension, as done by some existing approaches. However, we explore further by learning a regression model with this lower dimensional projection as the target, which allows us to predict the persona based on the user-context features. It will be presented in the next subsection.

B. Persona Prediction

Detecting the mere presence of multiple personas and also associating them with rating events is not sufficient for generating accurate personalized recommendations. It is important to identify which persona is generating the request at recommendation time in a video session. We propose to solve this problem by learning a mapping from user context to the *persona indicator* (i.e. the lower dimensional projection of each rating event as described in the last subsection).

$$f: \ U \times \mathcal{C}_3 \times \cdots \times \mathcal{C}_m \rightarrow \mathbb{R}^d.$$

Let the set of all context be C_{tx}, i.e., $\mathcal{C}_i \in C_{tx}$. We propose a non-parametric agglomerative clustering based approach to generate a non-linear regression model. It assumes the persona indicator values follow a conditional Gaussian distribution given specific circumstances specified by the context features. In other words, these values are Gaussian distributed conditioned on the joint state of user-context features.

The proposed approach is applied account-wise, which starts with smaller bins within an user account and iteratively combines these bins into larger bins if the defined metric is satisfied by the candidate small bins. Let set S_i^u denote all the context belonging to the ith session of an user u. Since, $d = 1$, let $f(U, S_i^u) = p_i$. In this case p_i is the first principal component denoted as 1PC value. Then, a unit bin B_i is defined by $B_i = \{p_i\}$. Unlike the standard agglomerative clustering which starts with each sample point as an independent bin, we start with the unit bins which are confined based on the joint state of the user-context features. Thus we expand B_i beyond the single element as:

$$S_k^u = S_i^u \implies p_k \in B_i.$$

The assumption behind this is that each persona has its specific user context (e.g. time, location, device, etc.). However, if multiple personas always share the common user context, it is impossible for any machine model to predict the persona solely based on the user context information.

Then, the algorithm iteratively looks at each one of the user-context features at a time and picks the candidate (to-be-combined) bins with different values of the current context feature given all the other context features being in the same joint state. Since we are only considering a single dimension of the projection, for combination metric, we choose the variance measure for decision of whether to combine multiple bins or not. If the variance of the 1PC values in the combined bin is smaller than all the variances from the candidate smaller bins plus a pre-defined *slack* variable ($\mathcal{S}$), then the smaller bins will be combined. We introduce the hyper-parameter 'slack' to relax the rigid bin combination metric and also to give us some control over how many bins are combined. It should be noted that the slack variable is specified as relative to the variance, and not in absolute terms. A few choices of this hyper-parameter can be $10\%, 25\%, 50\%$ etc. Then, the current context feature will be removed from the condition used to predict the 1PC value since it is uninformative w.r.t. the 1PC value. The resulting bins give the best description of the users' preferences under different circumstances specified by the most important context features.

$$
\begin{aligned}
&\{var(p_i)|p_i \in B_m \cup B_n\} \leq \\
&min[\{var(p_j)|p_j \in B_m\}, \{var(p_k)|p_k \in B_n\}] + \mathcal{S} \\
&\implies B_l = combine(B_m, B_n), S_l^u = S_m^u \cap S_n^u
\end{aligned}
$$

Since the results of the agglomerative clustering not only give the context which can be used to differentiate personas, but also conditional distributions of the 1PC value under every context. We used the mean value of the training sessions within each combined bin to estimate the 1PC value of the testing sessions which fall into the same bin. After we get all the predicted 1PC values for the testing sessions, we quantize these continuous values by employing Gaussian mixture models and selected the number of mixed components based on the BIC score. Those categorized 1PC values are finally added back into the input feature table as an additional column for training a factorization machine [10].

IV. Experiments

We incorporated our persona identification framework with a state-of-the-art recommender system, the Factorization Machines. We conducted three sets of experiments. First, we evaluate on how well our approach holds up under an extreme condition in which *every* account is shared. Next, we evaluate our approach on a dataset where the extent and effect of account sharing is unknown. Finally, we measure how accurate our identified personas are. For the first and second set of experiments we have compared our method with the FM model without our approach (we will refer to this as 'Vanilla-FM'), and also with FM trained on an account based clustering approach (referred as 'Cluster-FM'). The account based clustering approach clusters each individual account separately and considers each cluster as an atomic user. Some variants of the account based clustering approach is used by [3], [4], [8]. We determine the number of clusters using the silhouette score. For all of our experiments we have fixed the slack parameter to 25%, unless otherwise mentioned.

A. Quantitative Analysis: Simulated Data

We first evaluated our approach on a simulated dataset, where we know that each account is shared among multiple users. This dataset has around 25,000 rating events of 600 users on 198 videos, 1 user-context feature (i.e. Device) which has 2 possible values (Device A, Device B) and 1 video-context feature (i.e. Genre) which also has 2 possible values (Genre C, Genre D). The data has been generated as follows - each user prefers one genre (e.g., Genre C) from one device (e.g., Device A) and prefers the other genre (i.e., Genre D) when using the other device (Device B). The t-SNE visualization of the dataset is shown in Fig. 2. Here, blue dots indicate 'dislike' rankings while red dots indicate 'like' rankings. We further divided the dataset into training and test sets with a 90-10 split.

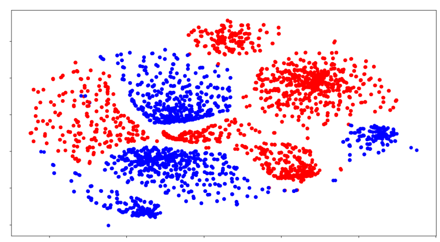

Fig. 2. TSNE visualization of simulated dataset.

On this dataset, we report the RMSE of the FM model with our approach along with the two competing methods on the test set shown in Fig 3(a). As we can clearly see, the RMSE of our approach is far superior compared to the competing two approaches.

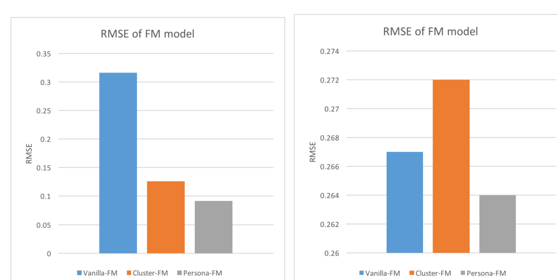

(a) Simulated (b) Industrial

Fig. 3. RMSE results on the two datasets

B. Quantitative Analysis: Industrial Data

We applied our approach to a real world dataset. The experiments were performed on the access logs for a period of one month. The dataset has 401,844 sessions of 12,022 users, 2005 videos, 9 different devices, 16 different genres, and 550 different geo-regions. There is no explicit rating available for the users and videos. Hence, we have used how much of a video has been watched (i.e.,'session progress') as a proxy for ratings (as outlined in [11], [12]). We selected 3000 sessions at random as a test set and the rest of the sessions as the training set.

The RMSE of our approach on the test dataset as compared to the other two approaches are shown in Fig 3(b). We can clearly see that our approach is slightly better than the other two approaches, however the improvements are less dramatic than the simulated dataset. It is also surprising that account based clustering is performing worse than the Vanilla FM. We believe that this is because when we split an account into multiple users there are far less data to train the 'Cluster-FM' faithfully. In this dataset (as with many other real world data from similar domains), the number of accounts with multiple personas is unknown and potentially low. However, the problem is sever for these accounts and hence need to be solved. However, we don't have the ground truth to compare only those accounts with multiple personas.

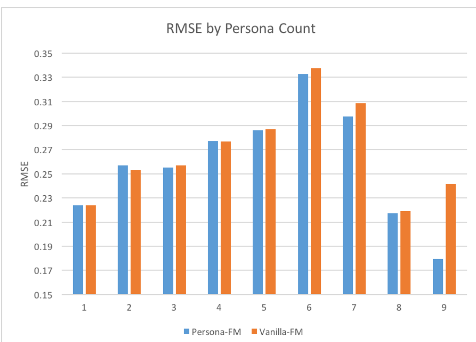

Fig. 4. Persona count wise RMSE results on the industrial data

Finally, we investigate the performance of our approach for different personas. Our non-parametric approach has chosen the maximum number of possible persona as 10. We examined the RMSE of our approach for the users with 1 persona, 2 persona etc. We then compared the RMSE for those users under Vanilla-FM. Our results are shown in Fig. 4. We observe that our approach performs almost as good as Vanilla-FM when an account has a few identified personas (1-5), whereas for users with a relatively large number of identified personas, our approach outperforms Vanilla-FM. This implies that those users might indeed have shared accounts (or at least distinct preference under different situations).

C. Cluster Performance: Industrial Data

We also wanted to verify the quality of our identified personas. However, without any ground truth it is difficult to *objectively* evaluate a clustering algorithm. Hence, in the absence of any external ground truth we decided to modify the real dataset in a controlled way as follows. We first identified accounts with fairly uniform viewing preference. Then we randomly sampled from them and fused pairs of accounts together. In this modified dataset we applied our approach and evaluated 1) how well separated the original accounts are and 2) how much over separation is in the original accounts.

Slack	Overlap	Purity
10%	0.208	0.688
25%	0.208	0.692
50%	0.186	0.688

TABLE I
OVERLAP AND PURITY UNDER DIFFERENT SLACK CRITERIA

For the first metric we define the measure overlap = (No. of personas with common between two account) / (Total no. of sessions). Ideally we would want an overlap of 0, which can be trivially achieved if each session is assigned to a different persona. And for the next metric we use a purity measure. Ideally we would want purity of 1. However that can be trivially achieved if all sessions of all the accounts have been assigned to the same persona. The result for the overlap and purity under different slack conditions are shown in Table I.

V. CONCLUSION

Generating accurate and personalized recommendation when there are multiple individuals sharing the same user account is a challenging problem, which can greatly affect performance of recommender systems. In this paper we propose a novel framework to not only detect the presence of account sharing but also to predict the user who is requesting a new recommendation. Our experiment on synthetic as well as real dataset has shown strength in our approach.

REFERENCES

[1] J. Masthoff, "Group Recommender Systems: Combining Individual Models," in *Recommender Systems Handbook*. Springer, 2011, pp. 677–702.
[2] Z. Yu, X. Zhou, Y. Hao, and J. Gu, "TV program recommendation for multiple viewers based on user profile merging," in *User Modeling and User-Adapted Interaction*. Springer, 2006, pp. 63–82.
[3] Z. Wang, Y. Yang, L. He, and J. Gu, "User Identification within a Shared Account: Improving IP-TV Recommender Performance," in *East European Conference on Advances in Databases and Information Systems*. Springer, 2014, pp. 219–233.
[4] Y. Yang, Q. Hu, L. He, M. Ni, and Z. Wang, "Adaptive Temporal Model for IPTV Recommendation," in *International Conference on Web-Age Information Management*. Springer, 2015, pp. 260–271.
[5] K. Iguchi, Y. Hijikata, and S. Nishida, "Individualizing user profile from viewing logs of several people for TV program recommendation," in *International Conference on Ubiquitous Information Management and Communication*, no. 61. ACM, 2015.
[6] P. G. Campos, A. Bellogn, I. Cantador, and F. Dez, "Time-Aware Evaluation of Methods for Identifying Active Household Members in Recommender Systems," in *Conference of the Spanish Association for Artificial Intelligence*. Springer, 2013, pp. 22–31.
[7] P. Bajaj and S. Sekhar, "Experience Individualization on Online TV Platforms through Persona-based Account Decomposition," in *ACM Multimedia Conference (MM)*. ACM, 2016, pp. 252–256.
[8] A. Zhang, N. Fawaz, S. Ioannidis, and A. Montanari, "Guess Who Rated This Movie: Identifying Users Through Subspace Clustering," in *Conference on Uncertainty in Artificial Intelligence (UAI)*. ACM, 2012, pp. 944–953.
[9] K. Verstrepen and B. Goethals, "Top-N Recommendation for Shared Accounts," in *Conference on Recommender Systems (RecSys)*. ACM, 2015, pp. 59–66.
[10] S. Rendle, "Factorization machines with libfm," *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 3, no. 3, p. 57, 2012.
[11] G. Wu, V. Swaminathan, S. Mitra, and R. Kumar, "Online video session progress prediction using low-rank matrix completion," in *Multimedia and Expo Workshops (ICMEW), 2014 IEEE International Conference on*. IEEE, 2014, pp. 1–6.
[12] G. Wu, V. Swaminathan, S. Mitra, and R. Kumar, "Context-aware video recommendation based on session progress prediction," in *Multimedia and Expo (ICME), 2017 IEEE International Conference on*. IEEE, 2017.