

Social Collaborative Filtering for Cold-start Recommendations

Suvash Sedhain*
ANU/NICTA and Kobo Inc.
Canberra, Australia
suvash.sedhain@anu.edu.au

Scott Sanner
NICTA/ANU
Canberra, Australia
ssanner@nicta.com.au

Darius Brazianus
Kobo Inc.
Toronto, Canada
dbraziunas@kobo.com

Lexing Xie
ANU/NICTA
Canberra, Australia
lexing.xie@anu.edu.au

Jordan Christensen
Kobo Inc.
Toronto, Canada
jchristensen@kobo.com

ABSTRACT

We examine the cold-start recommendation task in an online retail setting for users who have not yet purchased (or interacted in a meaningful way with) any available items but who have granted access to limited side information, such as basic demographic data (gender, age, location) or social network information (Facebook friends or page likes). We formalize neighborhood-based methods for cold-start collaborative filtering in a generalized matrix algebra framework that does *not* require purchase data for target users when their side information is available. In real-data experiments with 30,000 users who purchased 80,000+ books and had 9,000,000+ Facebook friends and 6,000,000+ page likes, we show that using Facebook page likes for cold-start recommendation yields up to a 3-fold improvement in mean average precision (mAP) and up to 6-fold improvements in Precision@k and Recall@k compared to most-popular-item, demographic, and Facebook friend cold-start recommenders. These results demonstrate the substantial predictive power of social network content, and its significant utility in a challenging problem – recommendation for cold-start users.

Categories and Subject Descriptors

H.4 [Information Search and Retrieval]: Information Filtering

General Terms

Algorithms; Experimentation; Performance

Keywords

Recommender systems; Cold-start problem

*Parts of this work were done while the first author was an intern at Kobo Inc.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or coresysercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. RecSys'14, October 6–10, 2014, Foster City, Silicon Valley, CA, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2668-1/14/10 ...\$15.00.

<http://dx.doi.org/10.1145/2645710.2645772>.

1. INTRODUCTION

The user cold-start problem concerns the task of recommending items to users who have not previously purchased or otherwise expressed meaningful preferences towards any items under consideration for recommendation. Addressing the cold-start problem can be important for first-time user engagement and retention and is therefore of critical significance in settings such as online retail.

While traditional collaborative filtering methods [1, 6] often provide strong recommendation performance once a user has made a few item purchases or preference judgments, they cannot be used in the cold-start setting when there are no user interactions with items. Content-based and user-based filtering [3] approaches may recommend solely based on user and item features in the absence of direct purchase or preference information over items; however, these methods typically provide very coarse approximations of content types (e.g., genre, author) and user profiles (e.g., gender, age, and location) and often only do marginally better than recommendation based solely on item popularity, as we show for the case of user-based demographic data in Section 4.

Recent years have seen the advent of social extensions of collaborative filtering [5] and efforts to leverage rich user information from Facebook and other social networks for predicting users' latent traits [2] and for recommendation [8, 10]. In the cold-start setting, Lin *et al* [4] leveraged social information for the *item* cold-start recommendation problem (where an item has not been previously purchased or rated by anyone), e.g., recommending new apps by leveraging Twitter feedback on an app's developer even though the app itself may be yet unrated.

Here, we propose to directly leverage user's social network content such as Facebook friends and page likes in a novel extension of item-based collaborative filtering formalized in a generalized matrix algebra framework. This framework does *not* require item purchase or preference data for target users and hence addresses the *user* cold-start problem.

We experiment with a subset of data from Kobo Inc. that contains 30,000 users who purchased 80,000+ ebooks and had 9,000,000+ Facebook friends and liked 6,000,000+ pages. We demonstrate that using Facebook page likes for cold-start recommendation yields up to a 3-fold improvement in mean average precision (mAP) and up to 6-fold improvements in Precision@k and Recall@k when compared to most-popular-item, demographic, and Facebook friend cold-

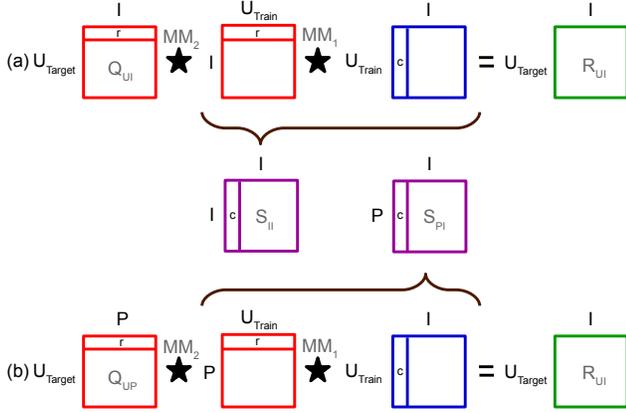


Figure 1: (a) A matrix algebra view of standard item-item neighborhood-based recommendation and (b) a variation for social cold-start recommendation that instead uses page-item similarity and does not require item purchase information for target users. U represents the user dimension (subdivided into train and test users), I the item dimension, and P the user’s personal information dimension (demographics, friends, or page likes). The \star operators annotated by MM_1 and MM_2 denote *generalized matrix multiplication, permitting any similarity metric over two vectors in place of the usual inner product.*

start recommendation approaches. These results demonstrate the substantial predictive power of social network content when leveraged in our novel social collaborative filtering framework for cold-start recommendation.

2. BACKGROUND AND CONTRIBUTION

We begin our discussion with a brief technical review of neighborhood-based collaborative filtering methods as viewed from a generalized matrix algebra framework and then discuss our extensions of this framework to a social collaborative filtering approach for cold-start recommendation.

2.1 Item-based Recommendation

We refer the reader to Figure 1(a) which shows a matrix view of item-based collaborative filtering [7]. Here Q_{UI} denotes a matrix of user-item purchases (users in rows, items in columns) for a set of target users $U_{Target} = \{u_1, u_2, \dots, u_m\}$ and a set of items $I = \{i_1, i_2, \dots, i_n\}$. For any $u \in U_{Target}$ and $i \in I$, we define

$$Q_{u,i} = \begin{cases} 1 & \text{if } u \text{ purchased } i \\ 0 & \text{otherwise} \end{cases}$$

For a set of training users U_{Train} (potentially overlapping with U_{Target}) and analogous definitions of the remaining matrices on the LHS of the equality in Figure 1(a), we can apply operator \star annotated by MM_1 (e.g. matrix multiplication) to obtain S_{II} , where for items $i \in I$ and $j \in I$, S_{ij} represents an item-item similarity between i and j . Finally, applying operator MM_2 to compute $Q_{UI} \star S_{II}$ yields a recommendation matrix R_{UI} where for each target user $u \in U_{Target}$ and item $i \in I$, R_{ui} represents a real-valued rating. Given R_{ui} and sorting items in descending order for each u yields a ranked list of most- to least-recommended items for user u .

By permitting definitions for the \star operators MM_1 and MM_2 other than standard matrix multiplication, we define a generalized matrix algebra framework for recommendation that can recover many existing frameworks when the operators are defined appropriately. For example, in order to recover a version of item-based collaborative filtering [7], one could use cosine similarity in place of the inner product for MM_1 . In general, we explore the following similarity metrics $Sim(\mathbf{r}, \mathbf{c})$ for the row \mathbf{r} and column \mathbf{c} vectors used to define \star in MM_1 and MM_2 (they need not be the same):

- **IP:** Standard matrix multiply inner product given by

$$Sim(\mathbf{r}, \mathbf{c}) = \langle \mathbf{r}, \mathbf{c} \rangle.$$

- **BinIP:** For $\tau \in \mathbb{R}$, a binary thresholded version of the inner product given by

$$Sim(\mathbf{r}, \mathbf{c}) = \begin{cases} 1 & \text{if } \langle \mathbf{r}, \mathbf{c} \rangle > \tau \\ 0 & \text{otherwise} \end{cases}.$$

- **LogIP:** A logarithm of the inner product

$$Sim(\mathbf{r}, \mathbf{c}) = \log \langle \mathbf{r}, \mathbf{c} \rangle.$$

- **Cos:** Cosine similarity which is obtained by an inner product of two L_2 normalized vectors equivalent to

$$Sim(\mathbf{r}, \mathbf{c}) = \frac{\langle \mathbf{r}, \mathbf{c} \rangle}{\|\mathbf{r}\| \|\mathbf{c}\|}.$$

While there are alternate similarity metrics that can be considered (e.g., Pearson Correlation Coefficient in general and Jaccard in the special case of binary vectors), the above metrics worked best for our cold-start evaluation.

2.2 Social Cold-start Recommendation

Returning to our discussion of Figure 1(a), since Q_{UI} requires item interaction information for the target users U_{Target} , item-based collaborative filtering is not applicable in the cold-start case. However, if we view the purchased subset of items I simply as user attributes and extend this attribute set to other non-item personal attributes P (e.g., demographics, Facebook friends, or page likes) then for users U_{Target} we can substitute P for I to obtain Figure 1(b).

More formally, we define $P = \{p_1, p_2, \dots, p_l\}$ as a set of attributes for user demographic traits (gender, age group, location), Facebook friends, or Facebook page likes. Then we define the matrices in Figure 1(b) analogous to 1(a) where the first two matrices replace dimensions over I with dimensions over P . Thus for Q_{UP} and any $p \in P$ (e.g., page likes) and $u \in U_{Target}$, $Q_{up} = 1$ if u had attribute p (e.g., u liked p) and $Q_{up} = 0$ otherwise. Defining the second matrix analogously, we apply $MM_1 \star$ to obtain S_{PI} which relates personal attributes P to preferences over items I . And finally $Q_{UP} \star S_{PI}$ yields recommendation matrix R_{UI} analogous to Figure 1(a) except that in 1(b), no interactions on I were required for the target recommendation users U_{Target} , thus making 1(b) applicable in the cold-start case.

While this recommender can use both social (e.g. friends, page likes) and non-social (e.g. demographics) information to define P , we note that our experimental evaluation in Section 4 shows that page likes provide an exceptionally strong signal for cold-start recommendation. These results lead us to focus on the contribution of this novel cold-start recommender in conjunction with social data.

k	Most Popular		Demographics		Friend Network		Page Likes	
	Precision	Recall	Precision	Recall	Precision	Recall	Precision	Recall
@1	0.006+/-0.003	0.006+/-0.003	0.014+/-0.004	0.014+/-0.004	0.014+/-0.006	0.014+/-0.006	0.038+/-0.015	0.038+/-0.015
@3	0.008+/-0.003	0.023+/-0.010	0.010+/-0.003	0.031+/-0.008	0.011+/-0.004	0.032+/-0.012	0.025+/-0.010	0.076+/-0.029
@5	0.008+/-0.004	0.039+/-0.019	0.010+/-0.004	0.050+/-0.018	0.009+/-0.003	0.044+/-0.016	0.023+/-0.009	0.117+/-0.044
@10	0.007+/-0.003	0.074+/-0.034	0.008+/-0.003	0.083+/-0.029	0.006+/-0.002	0.062+/-0.022	0.017+/-0.006	0.173+/-0.059
@20	0.006+/-0.003	0.120+/-0.067	0.006+/-0.003	0.125+/-0.051	0.004+/-0.001	0.078+/-0.025	0.010+/-0.003	0.205+/-0.063
mAP	0.026+/-0.011		0.035+/-0.010		0.029+/-0.010		0.075+/-0.025	

Table 1: Comparison of Most Popular baseline with various cold-start recommender systems.

k	IP-IP		LogIP-IP		BinIP-IP		Cos-IP	
	Precision	Recall	Precision	Recall	Precision	Recall	Precision	Recall
@1	0.027+/-0.013	0.027+/-0.013	0.030+/-0.012	0.030+/-0.012	0.028+/-0.012	0.028+/-0.012	0.030+/-0.012	0.030+/-0.012
@3	0.018+/-0.007	0.053+/-0.020	0.020+/-0.008	0.061+/-0.023	0.022+/-0.009	0.066+/-0.026	0.022+/-0.008	0.066+/-0.024
@5	0.017+/-0.006	0.086+/-0.032	0.019+/-0.008	0.097+/-0.038	0.020+/-0.008	0.102+/-0.040	0.020+/-0.008	0.102+/-0.038
@10	0.014+/-0.005	0.136+/-0.047	0.015+/-0.005	0.148+/-0.050	0.015+/-0.005	0.153+/-0.051	0.015+/-0.005	0.154+/-0.052
@20	0.009+/-0.003	0.181+/-0.060	0.009+/-0.003	0.185+/-0.061	0.009+/-0.003	0.186+/-0.060	0.010+/-0.003	0.191+/-0.060
mAP	0.057+/-0.019		0.062+/-0.021		0.063+/-0.022		0.065+/-0.021	

Table 2: Performance of Page Likes Recommender for $MM_1 \in \{\text{IP}, \text{LogIP}, \text{BinIP}, \text{Cos}\}$, $MM_2 = \text{IP}$.

k	IP-Cos		LogIP-Cos		BinIP-Cos		Cos-Cos	
	Precision	Recall	Precision	Recall	Precision	Recall	Precision	Recall
@1	0.031+/-0.014	0.031+/-0.014	0.031+/-0.013	0.031+/-0.013	0.030+/-0.014	0.030+/-0.014	0.038+/-0.015	0.038+/-0.015
@3	0.021+/-0.008	0.064+/-0.023	0.022+/-0.008	0.065+/-0.023	0.021+/-0.008	0.063+/-0.024	0.025+/-0.010	0.076+/-0.029
@5	0.020+/-0.008	0.100+/-0.039	0.020+/-0.008	0.101+/-0.038	0.020+/-0.007	0.100+/-0.037	0.023+/-0.009	0.117+/-0.044
@10	0.015+/-0.005	0.151+/-0.052	0.015+/-0.005	0.153+/-0.052	0.016+/-0.005	0.159+/-0.055	0.017+/-0.006	0.173+/-0.059
@20	0.010+/-0.003	0.193+/-0.061	0.010+/-0.003	0.193+/-0.061	0.010+/-0.003	0.202+/-0.063	0.010+/-0.003	0.205+/-0.063
mAP	0.065+/-0.022		0.065+/-0.022		0.065+/-0.021		0.075+/-0.025	

Table 3: Performance of Page Likes Recommender for $MM_1 \in \{\text{IP}, \text{LogIP}, \text{BinIP}, \text{Cos}\}$, $MM_2 = \text{Cos}$.

3. EXPERIMENTAL SETUP

The dataset we use in this study comes from Kobo Inc., a major online ebook retailer with more than 20 million readers. It contains an anonymized dataset of ebook purchases and Facebook friends and page likes for a random subset of 30,000 Kobo users; these users purchased more than 80,000 books, had over 9 million friends and liked about 6 million pages. We subsample pages by including only those pages liked by at least 5 people and not more than 5,000 people, which reduces the number of unique pages to about 600,000. The dataset also includes basic user demographics, namely age group, gender and location.

We split the dataset into 10 temporally divided train and test folds. To prepare the training dataset, we include all the user-item data prior to a specific date, starting from June 2012 and incremented by a month in each fold. The test set includes the first purchase of all new users in the following month when the user had at least 10 page likes.

We can obtain different social cold-start recommenders by choosing the similarity metrics (**IP**, **LogIP**, **BinIP**, **Cos**) used for both MM_1 and MM_2 in Figure 1(b). We express the choices in the order of MM_1 - MM_2 and show evaluations for the following eight possibilities (the remaining possibilities did not approach the best result reported among these eight):

- $MM_2 = \text{IP}$, $MM_1 \in \{\text{IP}, \text{LogIP}, \text{BinIP}, \text{Cos}\}$
- $MM_2 = \text{Cos}$, $MM_1 \in \{\text{IP}, \text{LogIP}, \text{BinIP}, \text{Cos}\}$

We used the threshold $\tau = 2$ for **BinIP**.

We compare the popularity-based baseline with three cold-start collaborative filtering recommender systems defined in Section 2.2, all using respective P as defined below and **Cos-Cos** similarities unless otherwise specified:

- **Most popular** baseline recommends the most popular items in the dataset in order of popularity. While this is a non-personalized recommender, it is often used in the cold-start setting.
- **Demographics** defines P as the set of binary attribute dimensions for a user’s gender, age-group and location. Here, we use 10 different disjoint age group indicators, each with a range of 10 years. Every unique location gets its own binary attribute.
- **Friends** defines P as the set of binary attribute dimensions for Facebook friends and assumes that all users are friends with themselves.
- **Page Likes** defines P as the set of binary attributes for Facebook page likes.

In all of the experiments, we report Precision@ k , Recall@ k and mean average precision (mAP@100)¹ using 10-fold cross validation and provide standard error bars corresponding to 95% confidence intervals.

4. RESULTS AND ANALYSIS

In Table 1, we compare our most popular baseline with three variants of our cold-start approach. Here, we see that social cold-start collaborative filtering using **Page Likes** yields up to a 3-fold improvement in mean average precision (mAP) and up to 6-fold improvements in Precision@ k and Recall@ k when comparing to **Most Popular**, **Demographic**, and Facebook **Friend Network** recommenders.

¹As defined in <https://www.kddcup2012.org/c/kddcup2012-track1/details/Evaluation> and used as a surrogate measure of area under the precision/recall curve.

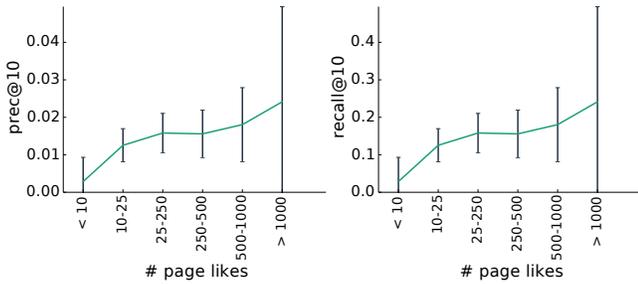


Figure 2: Performance of the Page Likes recommender for users vs. different numbers of page likes.

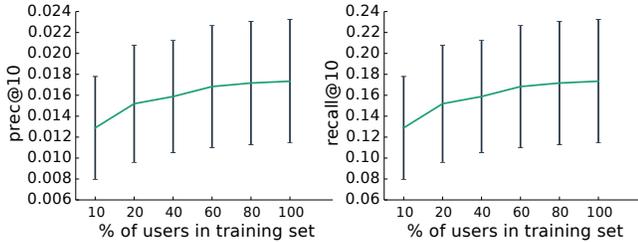


Figure 3: Performance of the Page Likes recommender vs. the percentage of users in U_{Train} used.

To assess whether **Cos-Cos** is indeed the best choice for MM_1 - MM_2 , we evaluate variants as previously described in Tables 2 and 3 and show that **Cos-Cos** is the best method for social cold-start recommendation with **Page Likes**. We remark that **Cos-Cos** likewise performed best among **Demographics** and the **Friend Network**.

To conclude our analysis we focus on the best performing **Page Likes** social cold-start recommender and investigate its performance as a function of different key quantities:

How does performance vary vs. the number of page likes for a target user? The number of page likes varies from user to user. Thus we divide users in U_{Target} into six categories based on the number of pages they have liked and evaluate the performance on each user category. Figure 4 shows that performance increases as the number of user page likes increases. However, the variance for users with the largest number of page likes is much higher, indicating that a lack of selectiveness with page likes can lead to high noise when using this information for cold-start recommendation.

How does performance vary vs. the number of users and pages in the training set? To analyze this for the number of users, we randomly select $x\%$ of users from U_{Train} and show the cold-start performance vs. x in Figure 4. This figure shows that the performance steadily increases with the number of users, but with diminishing returns beyond 80% of the users. Similarly, we randomly select $x\%$ of the dimensions in P and show the cold-start performance vs. x in Figure 4. This figure shows that the performance increases gradually with the number of pages but with less increase in improvement once 40% of the pages in P are used.

5. CONCLUSION

We defined a novel social collaborative filtering framework that generalizes standard item-based collaborative filtering

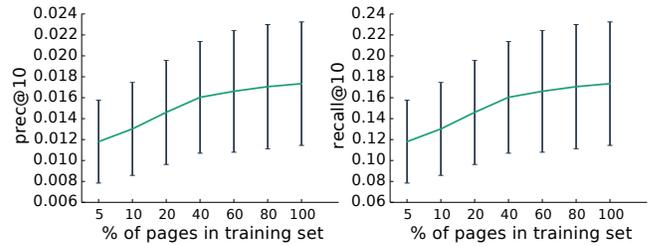


Figure 4: Performance of the Page Likes recommender vs. the percentage of pages in P used.

to the cold-start recommendation setting. When used in conjunction with Facebook page likes data for each user, this approach substantially outperformed cold-start recommenders based on popular items, demographics, and the Facebook friend network. Overall, these results demonstrate the substantial power of Facebook page likes when leveraged to address the user cold-start recommendation problem in a social collaborative filtering framework. Future work may explore whether this work can be extended with learning-based techniques such as *collective matrix factorization* [9].

6. ACKNOWLEDGMENTS

NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program. This research was supported in part by Australian Research Council project DP140102185.

7. REFERENCES

- [1] R. M. Bell and Y. Koren. Scalable collaborative filtering with jointly derived neighborhood interpolation weights. In *ICDM-07*, 2007.
- [2] M. Kosinski, D. Stillwell, and T. Graepel. Private traits and attributes are predictable from digital records of human behavior. In *PNAS-13*, volume 110, 2013.
- [3] K. Lang. NewsWeeder: Learning to filter netnews. In *ICML-95*, 1995.
- [4] J. Lin, K. Sugiyama, M.-Y. Kan, and T.-S. Chua. Addressing cold-start in app recommendation: Latent user models constructed from twitter followers. In *SIGIR-13*. 2013.
- [5] H. Ma, H. Yang, M. R. Lyu, and I. King. Sorec: Social recommendation using probabilistic matrix factorization. In *CIKM-08*, 2008.
- [6] P. Resnick and H. R. Varian. Recommender systems. *Communications of the ACM*, 40, March 1997.
- [7] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl. Item-based collaborative filtering recommendation algorithms. In *WWW-01*, 2001.
- [8] S. Sedhain, S. Sanner, L. Xie, R. Kidd, K.-N. Tran, and P. Christen. Social affinity filtering: Recommendation through fine-grained analysis of user interactions and activities. In *COSN-13*, 2013.
- [9] A. P. Singh and G. J. Gordon. Relational learning via collective matrix factorization. In *KDD-08*, 2008.
- [10] Y. Zhang and M. Pennacchiotti. Predicting purchase behaviors from social media. In *WWW-13*, 2013.