

For more papers on (research) paper recommender systems, visit www.docear.org

Persistence in Recommender Systems: Giving the Same Recommendations to the Same Users Multiple Times

Joeran Beel¹, Stefan Langer¹, Marcel Genzmehr¹, and Andreas Nürnberger²

¹Docear, Germany

{beel|langer|genzmehr}@docear.org

²Otto-von-Guericke University, Dpt. of Computer Science, DKE Group, Magdeburg, Germany

andreas.nuernberger@ovgu.de

Abstract. How do click-through rates vary between research paper recommendations previously shown to the same users and recommendations shown for the very first time? To answer this question we analyzed 31,942 research paper recommendations given to 1,155 students and researchers with the literature management software Docear. Results indicate that recommendations should only be given once. Click-through rates for ‘fresh’, i.e. previously unknown, recommendations are twice as high as for already known recommendations. Results also show that some users are ‘oblivious’. It frequently happened that users clicked on recommendations they already knew. In one case the same recommendation was shown six times to the same user and the user clicked on it each time again. Overall, around 50% of clicks on reshown recommendations were such ‘oblivious-clicks’.

Keywords: recommender systems, persistence, re-rating, research paper

1 Introduction

Recommender systems became popular in many domains during the past decades and content-based and collaborative filtering became the two most dominant approaches. Some researchers in the field of collaborative filtering analyzed the effect of letting users re-rate items. They found that correlation between original ratings and new ratings was low and only 60% of users gave the same rating as before [1]. Amatriain et al. showed that it might be better to letting users re-rate items than showing new ones. By doing so accuracy of recommender systems increased by around 5% [2].

We wonder whether re-showing recommendations might make sense in general. For instance, a user might miss a recommendation the first time, simply because he was in a hurry and did not pay attention to the recommendation. In this case it would make sense for a recommender to be persistent and to display the same recommendation again. To the best of our knowledge ‘recommendation persistence’ has not been studied so far.

2 Research Objective & Methodology

Our goal was to find out if and how often it makes sense to display the same recommendations to the same users. To answer this question we analyzed empirical data from the literature management software Docear [4] which features a research paper recommender system [3]. The recommender system recommends research papers to users regardless of whether papers were previously recommended to the users or not. We analyzed how

click-through rates (CTR) between recommendations shown only once and CTR of recommendations shown multiple times differed. CTR expresses how much percent of the delivered recommendations were clicked. For instance, if 12 recommendations were clicked out of 1,000 delivered ones, CTR would be 1.2%. CTR basically measures the ‘precision’ of the recommendation algorithm under the assumption that a clicked recommendation is a ‘good’, i.e. useful, recommendation. For further details on Docear and its recommender system (e.g. how recommendations are generated and displayed) see [3, 4].

3 Results

31,942 recommendations were shown to 1,155 users for the first time and from the 31,942 recommendations 1,677 were clicked, which equals a click-through rate of 5.25% (Table 1). From the 31,942 recommendations 2,466 were shown a second time to 375 distinct users and 154 recommendations were clicked (CTR 6.24%). From the 2,466 recommendations 574 were displayed a third time and CTR was 6.97%. Also for the fourth iteration CTR was still rather high (6.55%). Based on these results one might conclude that it could make sense to display recommendations at least two or three times because for these reiterations CTR was significantly higher than for the first one ($p < 0.05$).

Table 1: Reiterations and click-through rate (CTR)

	Reiteration									
	1	2	3	4	5	6	...	11	...	21
Users	1,155	375	97	38	12	6		-		1
Impressions	31,942	2,466	574	229	112	71		2		1
No clicks	30,265	2,312	534	214	100	68		2		1
Clicks	1,677	154	40	15	12	3		-		-
CTR, overall	5.25%	6.24%	6.97%	6.55%	10.71%	4.23%		0.00%		0.00%
Obliv.-clicks	1st click	1,677	97	14	8	7	-	-		-
	2nd click	-	57	13	1	2	1	-		-
	3rd click	-	-	13	3	2	1	-		-
	4th click	-	-	-	3	-	-	-		-
	5th click	-	-	-	-	1	-	-		-
	6th click	-	-	-	-	-	1	-		-
	Σ Obliv. clicks	-	57	26	7	5	3	-		-
% Obliv. clicks	0%	37%	65%	47%	42%	100%		-		-
CTR, 1st click	5.25%	3.93%	2.44%	3.49%	6.25%	0.00%		0.00%		0.00%

The picture changes when looking at more detail into the data: around 50% of all clicks on reshown recommendations are ‘oblivious-clicks’ (Table 1, lower part). We define an ‘oblivious click’ as a click on a recommendation that the user should know already, because he clicked it previously. For instance, 574 recommendations were shown three times. 40 of these recommendations were clicked which equals a CTR of 6.97%. However, only 14 were clicked for the first time – the other 26 (2x13) were clicked for the second or even third time. In one case a recommendation was even shown six times to the same user and the user clicked it each time. Ignoring the oblivious-clicks, i.e. considering only 1st clicks, CTR decreases the more often recommendations are shown. Therefore, results may indicate that CTR increases when showing recommendations multiple times but only because users sometimes clicked on recommendations they have clicked before.

In addition, CTR increased in general the more recommendations were shown previously to a user (Figure 1). For instance, CTR did not only increase for reshown recom-

mentations but also for ‘fresh’ recommendations, i.e. recommendations being displayed to a user for the very first time. This is not surprising because users who receive many recommendations probably are using the software for a longer time than users receiving their first recommendations. And for users using the software for a longer time, better user models can be created and hence better recommendations can be given (although this is not always the case as shown in [5].

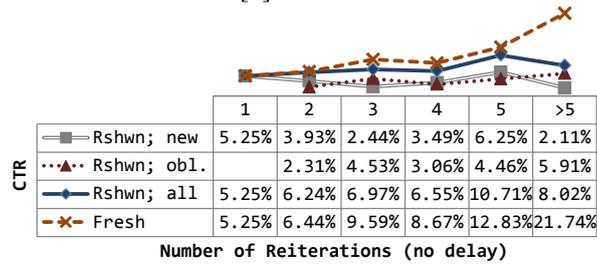


Figure 1: Redisplayed recommendations vs. fresh ones

To get a better understanding of how good re-shown recommendations performed, we compared their CTR with CTR of fresh recommendations. If a recommendation was shown the second time, it received a CTR of 6.24% on average – a CTR of 3.93% for reshown recommendations not being clicked before and a CTR of 2.31% for reshown recommendations being clicked before (Figure 1). In contrast, fresh recommendations being displayed at the same time achieved a CTR of 6.44% and hence performed better than the reshown recommendations. This is true for all iterations: fresh recommendations always performed better than reshown recommendations at the same time (including oblivious-clicks). Considering only new clicks on reshown recommendations (i.e. ignoring oblivious clicks), fresh recommendations performed even two to three times as good.

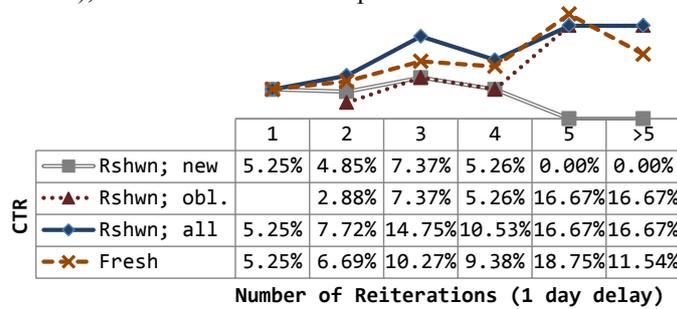


Figure 2: Fresh recommendations vs. redisplayed ones with at least one day delay

Based on the presented numbers one could conclude that reshown recommendations would never make sense. However, we did the same analysis for recommendations that were reshown with at least one day delay (Figure 2). That means we ignored all recommendations in the analysis that were reshown to the same user within 24 hours. In this case, CTR of reshown recommendations is often better than for fresh recommendations (with oblivious-clicks included). For instance, for the second iteration CTR for fresh recommendations was 6.69% but for reshown recommendations 7.72%. However, when ignoring oblivious-clicks again fresh recommendations always perform better than reshown recommendations. We also conducted the same analysis with a longer delay (three, seven, and fourteen days). Results were similar to the ones presented. Due to space restrictions we omit further details.

4 Interpretation and Outlook

Our results indicate that it makes no sense to generally display recommendations multiple times to the same users – fresh recommendations usually perform better. Nevertheless, about 2-3 % of recommendations shown the second or third time were clicked by the users for the first time. By showing recommendations only once, researchers would miss this 2-3% of interesting articles. In further research it should be studied why users sometimes click recommendations only when they were shown multiple times and whether users eventually found those recommendations useful or not. If they found the recommendations useful, then it should be studied how to find out which recommendations to show multiple times and how often. For instance, it might be that the interest of a user has changed – maybe even due to the recommendations he has seen – and on first display the recommendation simply was not relevant for him. That means if a strong concept drift was determined by the recommender system, recommendations shown previously (before the concept drift) might be given again.

In addition, it should be studied why users click several times on the same recommendations. We assumed that users were just oblivious. In this case it probably would be of little benefit for the user to see the same recommendations several times. But maybe obliviousness is not the only reason for clicking recommendations multiple times.

It is also quite interesting that it made a difference whether a recommendation was re-shown before or after 24 hours of a previous impression. In latter case (delay of one day or more), click through rates were significantly higher than for recommendations being re-shown within 24 hours and CTR of the reshown recommendations was even higher than for fresh recommendations. Under the assumption that oblivious clicks are desirable, reshowing recommendations could make sense. It might also make sense to transfer this finding to collaborative filtering and study how long to set a delay before letting users re-rate their items.

Open Data. Due to space restrictions, some data and graphs were omitted in this paper. For those being interested in more details (or validating our research), we publish our data on <http://labs.docear.org>.

5 References

1. Cosley, D., Lam, S.K., Albert, I., Konstan, J.A., Riedl, J.: Is seeing believing?: how recommender system interfaces affect users' opinions. Proceedings of the SIGCHI conference on Human factors in computing systems. pp. 585–592. ACM (2003).
2. Amatriain, X., Pujol, J.M., Tintarev, N., Oliver, N.: Rate it again: increasing recommendation accuracy by user re-rating. Proceedings of the third ACM conference on Recommender systems. pp. 173–180. ACM (2009).
3. Beel, J., Langer, S., Genzmehr, M., Nürnberger, A.: Introducing Docear's Research Paper Recommender System. Proceedings of the JCDL (2013).
4. Beel, J., Gipp, B., Langer, S., Genzmehr, M.: Docear: An Academic Literature Suite for Searching, Organizing and Creating Academic Literature. Proceedings of the 11th international joint conference on digital libraries. pp. 465–466. ACM (2011).
5. Beel, J., Langer, S., Nürnberger, A., Genzmehr, M.: The impact of users' demographics (age and gender) and other characteristics on evaluating recommender systems. Proceedings of the (TPDL) (2013).