Research Repository UCD

Title	Power to the people : exploring neighbourhood formations in social recommender systems
Authors(s)	Bourke, Steven, McCarthy, Kevin, Smyth, Barry
Publication date	2011-09
Publication information	Bourke, Steven, Kevin McCarthy, and Barry Smyth. "Power to the People : Exploring Neighbourhood Formations in Social Recommender Systems." ACM, 2011.
Conference details	Paper presented at RecSys-11, Chicago IL, USA, October 23-27, 2011
Publisher	ACM
Item record/more information	http://hdl.handle.net/10197/3508
Publisher's statement	This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in the RecSys '11 Proceedings of the fifth ACM conference on Recommender systems http://doi.acm.org/10.1145/2043932.2043997
Publisher's version (DOI)	10.1145/2043932.2043997

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Power to the People: Exploring Neighbourhood Formations in Social Recommender Systems

Steven Bourke, Kevin McCarthy & Barry Smyth
CLARITY: Centre for Sensor Web Technologies
School of Computer Science and Informatics
University College Dublin Ireland
[firstname.lastname]@ucd.ie

ABSTRACT

The explosive growth in online social networks in recent times has presented a powerful source of information to be utilised in personalised recommendations, unsurprisingly there has already been a large body of work completed in the recommender system field to incorporate this social information into the recommendation process. In this paper we examine the practice of leveraging a user's social graph in order to generate recommendations. Using various neighbourhood selection strategies, we examine the user satisfaction and perceived trust in recommendations received through a Facebook application.

General Terms

Recommender systems, social recommender systems, user study, neighbourhood formation

Keywords

Recommender Systems, Social web, Social recommender systems

1. INTRODUCTION

When recommender systems were first proposed, very few people would have imagined the connected world in which we live today. Work in collaborative recommender systems focussed on algorithmic designs to help build neighbourhoods of similar users to inform the recommendation process. At the time the relationships between the users that make up these neighbourhoods was not considered important, however with the explosion in social content systems where users are connected to each other through a social graph means that these relationships cannot be ignored. In this paper, we look at how users go about manually forming neighbourhoods and how these neighbourhoods inform recommendations when compared to other neighbourhood selection strategies. We examine how a user's interpretation of the effectiveness of a neighbourhood is affected when the

neighbourhood is revealed to them and also the relationship between user's satisfaction rating of recommendations compared to that user's perceived trust in the source of those recommendations.

In the next section we will introduce some related work around the area of neighbourhood formations and social graphs. In Section 3, we will introduce our neighbourhood selection strategies and in Section 4 we evaluate the performance of our neighbourhood strategies through a live-user study of Facebook users. We also analyse the user's satisfaction and perceived trust in the received recommendations.

2. BACKGROUND

Collaborative filtering has proved to be a very powerful approach to finding items of interest to users in large scale systems. In collaborative filtering a key step of the recommendation process is the neighbourhood formation stage. The neighbourhood formation stage involves the clustering of similar users together to find possible items which can be recommended to the target user. The neighbourhood formation stage has had a large body of work done over the years which either revolved around the process of identifying possible users for neighbourhood formation by using different similarity metrics or different ways in which these neighbourhoods are formed such as thresholding or nearest neighbour[2]. However it is not always the case that these approaches make sense when using a recommender system despite their accurate performances in offline evaluations [9]. More recently we've seen services with underlying social networks become more prominent in peoples day to day lives. This has allowed for more meaningful interpretation of relationships between users to be incorporated into the recommendation process. Before social graph information was easily available, a number of alternative ways to forming user neighbourhoods were proposed. One technique used demographic information about the users when trying to generate recommendations [8]. The authors proposed using demographic information to find regularity amongst the types of users that liked particular items. Overall they found no significant performance increase, but did find it beneficial when used alongside a content based approach. Trust has also been a measure that researchers have used in the recommendation process [5]. Massa found an overall reduction in the rate of incorrect predictions through the use of trust. Additionally work which involved explanations of neighbourhood formations has also been carried out. In one piece of work the authors provided a visual tool to allow a target user to choose who their peers were

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 commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

RecSys 2011

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.

when finding recommendations [7]. However the peers in this case were not based on social graph information, therefore a users level of knowledge concerning their peers would be somewhat limited. However another approach explored explaining recommendations from social graph information [1]. The authors of this work demonstrate an approach to social recommendations, a familiarity metric was proposed where by a user interacting with the system would be shown a potential 'friends' score towards the same item they had just rated. The authors do conclude that people are likely to prefer recommendations from people they do know, and as such it could prove beneficial to include social data into recommendation systems. Similar conclusions were drawn from the work carried out by Lee[3], while work carried out from Bonhard[1] focussed more on an empirical study, the Lee study performed analysis on a online social network. Lee found that users who had a bi-directional relationship had far higher levels of similarity than unidirectional. The authors did propose challenges related to the dissemination of possible recommendations due to the recommendations only been circulated inside their social network. Ma[4] explores the integration of social graph information into the directly into the recommendation process. They make the key distinction that a trust based approach [5] will typically involve trusted users who the target user does not know in any social capacity. The argument is that typically these users may not have a strong overlap of similar items, where as users from a social graph will have a stronger item based similarity.

3. NEIGHBOURHOOD FORMATION

In this section we will describe the neighbourhood formation approaches that we used to inform our recommendation approach. We introduce four neighbourhood selection techniques, three of these techniques consider the users social graph when deciding neighbourhood formation, while the final technique calculates a purely algorithmic neighbourhood as per a standard collaborative filtering system.

S1 - Social Graph: User Selected In this technique the target user is provided with a means to manually select users from their own social graph to add to their recommendation neighbourhood. In this way the choice of neighbourhood is fully controlled by the user, presumably with members being selected based on the user's own relationship with these other users.

S2 - Social Graph: Communication Frequency In this technique the neighbourhood is once again formed from a subset of the user's social graph. However, instead of allowing the user to manually select users to add to the neighbourhood a simple communication frequency metric is used to automatically select those users who communicate most frequently with the target user. We calculate this based on a simple count of the interactions a user has had with our target user.

S3 - Social Graph: Similarity-Based This technique is similar to S2 except that users are selected based on the similarity to the target user. Specifically we use the Extended Jaccard similarity metric that compares two users based on their item features. This, the user's neighbourhood is composed of people they know from

the social graph who are most similar to them based on their learned preferences.

S4 - Global Similarity Based This final technique corresponds to the neighbourhood formation technique most commonly used in standard collaborative filtering technologies, meaning we use the wider community of system users going far beyond the users social graph. Specifically, the target user's neighbourhood is formed from the k most similar users by comparing users based on their item features. The essential point here is that it is unlikely that the target user will have any explicit social relationship with their neighbours; neighbours are unlikely to be present in the target user's social graph.

We wish to examine the effects each of these neighbourhood selection strategies has on recommendation satisfaction. In typical systems the neighbourhood is never revealed to the end user - this is because revealing the neighbourhood may not make sense to the user as it most likely will be made up of other users whom they do not recognise or know. However, neighbourhoods formed using our strategies S1, S2 and S3 will be made up of other users known to our user, so revealing neighbourhoods may have an impact on recommendation satisfaction. Another hypothesis is that recommendations which are generated using the preferences of close friends in a user's social graph will benefit from an increase in perceived trust i.e. users will trust recommendations generated using information from their friends over recommendations from strangers. We test the neighbourhood selection strategies, revealing of neighbourhoods and the perceived user trust in recommendations in the following section.

4. EVALUATION

To evaluate our work we carried out a live user study. In total, 82 participants took part and the trial ran over April/May 2011. The user study was implemented as a Facebook application which recommended movie/TV content. We used Facebook because it allowed us access to a user's social graph which would usually have an extensive list of preferences associated with it for each user. We could also extract the users friends, their profile picture and name. We felt that overall the level of familiarity a user would have with their friends on Facebook is likely to be higher than on other services.

When participants sign into our experimental application we extract the user's preferences as well as their friends interests which are explicitly stated in their Facebook profiles. From our experiment, the average number of movie and TV content preferences per profile was 30 and the average size of a user's social graph was 189 users.

The first step of the experiment asked the user to manually select friends they felt would be a good source from which to generate recommendations (see Figure 1). The user then proceeded to select their neighbourhood for use in the S1 neighbourhood selection strategy. The number of friends they selected set the default neighbourhood size for the other neighbourhood selection strategies (a minimum requirement for at least ten friend selections was enforced). Once the neighbourhood selection was complete the application calculated recommendations using the various neighbourhood selection strategies.

Figure 1: Friend selection process.

The preference information extracted from the Facebook profiles is unary in nature. This means that a user will either have an item associated with their user profile (through a “like” preference) or they will not. To produce recommendations, we use the technique proposed by Mild [6]. This approach uses the Extended Jaccard metric to determine user-to-user similarity, this is used because approaches that use Pearsons as a similarity metric will have a item rating which is used in the similarity calculations. The prediction algorithm is calculated by the similarity weights between the active user and each of the users neighbours who also have the item. Additionally we use inverse user frequency when calculating user similarities to penalise popular items.

Recommendations lists were then produced using each of the neighbourhood selection techniques ($S1-S4$). The recommendation lists were generated by selecting the top recommendation from each of the neighbourhood selection strategies in a random order. This produced a list of 4 recommendations (If a recommendation item was already assigned to the list the next item on the top N list was inserted). This process was continued until a list of 20 TV/movie recommendations were complete.

Before displaying the final recommendation list, each participant was randomly assigned to one of two setups. In the first setup (see top of Figure 2) recommendation layout consisted purely of the recommendation and the rating scale. In the second setup (see bottom of Figure 2) the recommendation layout was the same except that the neighbourhood which was utilized in the formation of that recommendation was revealed to the participant. In our evaluation, 43 participants received the default recommendation list (without neighbourhoods shown) where as 39 participants received the recommendation lists with the neighbourhoods revealed.

The recommendations were then displayed to the participant and they were asked to rate each of their 20 TV/Movie items on a typical 5 star rating scale. The final part of the experiment asked each participant to rate each of the neighbourhood selection strategies in terms of how much they would trust recommendations generated using those strategies. The user was shown pictures and names of the neighbours from each neighbourhood.

5. RESULTS

Our experimental results show that on average users would select 14 neighbours, the most popular selection size was 11

Figure 2: Both Recommendation layouts used in evaluation, the second layout reveals the neighbourhood used.

and the maximum neighbourhood size was 32. We found that on average users spent three minutes and thirty seven seconds selecting their potential neighbours.

Did this time investment in manually selecting a neighbourhood produce better quality recommendations? Also, how did the ratings of recommendations differ across the various neighbourhood selection strategies? The results are shown in Figure 3. The bars represent the average rating participants gave to recommendations across the different neighbourhood formations when the neighbourhood was hidden and revealed to the user. We can see that the average rating given to recommendations does not vary significantly across the neighbourhood formation strategies. This holds true for both recommendation list layouts, with participants rating recommendations approximately 3 for the layout without neighbourhoods shown, up to an average rating of approximately 3 and a half for the layout with the neighbourhoods revealed.

The extra time given over to manual neighbourhood suggestion does not produce much difference in the actual ratings participants give to recommendations no matter what neighbourhood formation strategy is used. However, the revealing of the neighbourhood to the participant does have an impact. When neighbourhoods were revealed participants rated these recommendations higher than when they had no knowledge of the neighbourhood. Surprisingly, even when participants could see the neighbourhoods, and more importantly, the neighbourhood they manually selected themselves, they still rated recommendations from all neighbourhood strategies in a similar manner.

Participants may not be as good at selecting useful friends

Figure 3: Average participant rating scores with the overall perceived trust overlay.

for neighbourhoods as they think they are. However it is worth noting that on average the computationally predicted score for recommendations from S1 was the lowest of the four strategies, with S4 getting the highest average predicted score. In conclusion, revealing the neighbourhoods used in the generation of recommendations does have an impact on how users rate recommendations however the formation and makeup of these neighbourhoods has little effect.

Regarding trust, our basic experimental hypothesis is that users will place more trust in recommendations which originate from neighbourhoods that they have a hand in selecting or that are influenced directly by their own social graph. If this hypothesis holds then we expect to see a significant difference in the average trust scores assigned to S1 (user defined neighbourhood) compared to S4 (purely algorithmic neighbourhood). It is not clear, however, the extent to which users will trust S2 and S3 over S4, if at all, or whether there will be a difference in the trust scores assigned to S2 (social graph neighbourhood based on frequency of contact/communication) versus S3 (social graph neighbourhood based on an independent model of inter-user similarity).

The average trust scores for each of the neighbourhood strategies is shown as line in Figure 3. We found that throughout the user study users gave high levels of perceived trustworthiness to neighbourhoods which were derived from their own social interactions as opposed to neighbourhoods which were algorithmically generated. The trust ratings remained the same for the groups of users who had the neighbourhoods revealed and those that did not. What is quite interesting is that S3 gets a much lower perceived trust score, despite the recommendations coming from the users social graph. This could be that comparatively, the user would trust their own pre-selected neighbours to a high regard, and due to familiarity of consistent communication will much prefer the neighbours from S2 over the neighbours from S3. As expected, users did not give high trust scores to S4.

6. CONCLUSION

In this paper we presented a number of different neighbourhood formation strategies and evaluated them in a live user study. There is not much difference in the actual ratings users give recommendations no matter what neighbourhood selection strategy is used even though the user perception is that they would trust recommendations powered by their selected neighbourhood/close friends more than they would from the wider community. There are improvements in recommendation ratings when the neighbourhoods are shown, however the trust perceptions remain the same. It is interesting to note that S1 and S2 gain the biggest performance increase when revealing where recommendations come from. Our future work will involve exploration of the social ties when generating recommendations for users, as well as effective ways in explaining recommendations from a social point of view.

7. ACKNOWLEDGEMENTS

This material is based on works supported by Science Foundation Ireland under Grant No. 07/CE/11147 CLARITY CSET

8. REFERENCES

- [1] Philip. Bonhard and M. Angela Sasse. Knowing me, knowing you - Using profiles and social networking to improve recommender systems. *BT Technology Journal*, 24(3):84–98, July 2006.
- [2] Jon Herlocker, Joseph A. Konstan, and John Riedl. An Empirical Analysis of Design Choices in Neighborhood-Based Collaborative Filtering Algorithms. In *Information Retrieval*, volume 5, pages 287–310–310. Springer Netherlands, 2002.
- [3] Danielle H. Lee and Peter Brusilovsky. *Social networks and interest similarity*. ACM Press, New York, New York, USA, June 2010.
- [4] Hao Ma, Dengyong Zhou, Chao Liu, Michael R. Lyu, and Irwin King. *Recommender systems with social regularization*. ACM Press, New York, New York, USA, February 2011.
- [5] Paolo Massa and Paolo Avesani. Trust-Aware Collaborative Filtering for Recommender Systems. volume 3290 of *Lecture Notes in Computer Science*, pages 492–508–508, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.
- [6] Andreas Mild. An improved collaborative filtering approach for predicting cross-category purchases based on binary market basket data. *Journal of Retailing and Consumer Services*, 10(3):123–133, May 2003.
- [7] John O'Donovan, Brynjar Gretarsson, Svetlin Bostandjiev, Tobias Hollerer, and Barry Smyth. *A Visual Interface for Social Information Filtering*. IEEE, August 2009.
- [8] Michael J. Pazzani. A Framework for Collaborative, Content-Based and Demographic Filtering. In *Artificial Intelligence Review*, volume 13, pages 393–408–408. Springer Netherlands, 1999.
- [9] Rachael Rafter, Michael O'Mahony, Neil Hurley, and Barry Smyth. *What Have the Neighbours Ever Done for Us? A Collaborative Filtering Perspective*, volume 5535 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.