A Constrained Spreading Activation Approach to Collaborative Filtering

Josephine Griffith[1], Colm O'Riordan[1], and Humphrey Sorensen[2]

[1] Dept. of Information Technology,
National University of Ireland, Galway, Ireland
josephine.griffith@nuigalway.ie, colm.oriordan@nuigalway.ie
http://www.it.nuigalway.ie/cirg/
[2] Dept. of Computer Science, University College Cork, Cork, Ireland
h.sorensen@cs.ucc.ie

Abstract. In this paper, we describe a collaborative filtering approach that uses features of users and items to better represent the problem space and to provide better recommendations to users. Features of the collaborative filtering dataset are found and incorporated into a network representation of the collaborative filtering space where users and items are represented by nodes and where the nodes are connected by weighted edges. A spreading activation approach to collaborative filtering, using this representation and constrained by the user and item feature information, is compared with a traditional collaborative filtering approach.

1 Introduction

There is much evidence from various domains within Information Retrieval that the combination of sources of evidence leads to more effective retrieval [8]. With respect to recommender systems, Basu et al. claim that "there are many factors which may influence a person in making choices, and ideally, one would like to model as many of these factors as possible in a recommendation system" [3]. Although the types of information combined and the techniques used vary substantially across the domains, it has been shown that there are advantages to be gained from considering more than one source of information. This paper considers additional information that can be obtained from the collaborative filtering data set itself and shows how, via a graph representation and a spreading activation approach, the incorporation of this information can aid in the collaborative filtering task.

The collaborative filtering problem space is often viewed as a matrix consisting of the ratings given by each user for items in a collection. Using this matrix, the aim of collaborative filtering is to predict the ratings of a particular user, i, for one or more items previously not rated by that user. The problem space can equivalently be viewed as a network where users and items are represented by nodes and the links between user nodes and item nodes represents a rating for the item by the user. Several researchers have adopted graph representations in

order to develop recommendation algorithms and a variety of graphs have been used and a number of graph algorithm approaches have been adopted [1], [11].

In this work, a graph model and spreading activation approach is used for recommendation. Certain features of the dataset are considered and, based on these features, certain nodes (or users and items) of the graph are identified prior to retrieval. These nodes are used to constrain the spreading activation approach thus resulting in only some of the edges being considered and only portions of the graph being traversed.

The paper outline is as follows: Section 2 presents related work; Section 3 outlines the approach taken; Section 4 details the experiments performed. Results are presented in Section 5 and Section 6 discusses conclusions and future work.

2 Related Work

Collaborative filtering techniques produce recommendations for some active user using the ratings of other users, where these users have similar preferences to the active user. It differs to traditional retrieval and filtering systems which return items to some user based on a comparison between the content contained in items (documents) and the content of a user query (information need). Given a set of users, a set of items, and a set of ratings, collaborative filtering systems attempt to recommend items based on the prior ratings of the users. The collaborative filtering system essentially automates the "word of mouth" process [19].

In various situations, collaborative filtering, on its own, is not adequate, e.g. when no ratings exist for an item; when new users join a system; or when a user is not similar to any other users in the dataset. Many different models and approaches have been proposed to overcome these problems. One proposed solution involves utilising any available additional information or evidence (collaborative, content, link, etc.).

Within the information retrieval and filtering domains, approaches have been developed for combining different types of information and combining results from different information retrieval systems [8]. Several authors suggest methods for combining content with collaborative information [2], [3], [15], [20]. There has also been some work on extracting additional information (or features) from the collaborative filtering data set and using this information to aid in recommendation [16], [17].

2.1 Graph-Based Approaches for Recommendation

Several researchers have adopted graph representations in order to develop recommendation algorithms [18], [14]. A variety of graphs have been used (e.g. directed and two-layer) and a number of graph algorithm approaches have been adopted. Aggarwal et al. present *horting*, a graph-based technique where nodes represent users and directed edges between nodes correspond to the notion of predictability [1]. Predictions are produced by traversing the graph to nearby

nodes and combining the ratings of the nearby users. Huang et al. present a number of graph-based representations of the collaborative filtering space using binary ratings. They test a number of associated algorithms, including spreading activation and link prediction approaches [11], [12].

2.2 Spreading Activation

The idea of spreading activation originated from the field of Psychology and was first used in Computer Science in the area of Artificial Intelligence to process semantic networks. Spreading activation approaches have been used in many Information Retrieval applications [5], [7] and more recently in the domain of collaborative filtering [11]. Spreading activation approaches have also been used to integrate sources of evidence and information [21], [6].

The spreading activation model involves a number of iterations where each iteration consists of one or more “hops” or “pulses” where a hop involves the spreading of activation from one node to all other nodes connected to it. A hop may also include a pre-adjustment stage and a post-adjustment stage which usually involves a decay factor. The activation of a node is calculated based on the inputs to the node. The input is found by summing the output from each of the nodes connected to it by the weight on the link connecting the two nodes. The output or activation of a node is usually a function of the input value where many different functions can be used (e.g., a threshold function). A number of problems with this basic spreading activation approach have been identified, one of which is that activation can spread to all nodes in just a few hops. This problem can be overcome by use of constraints in the pre- and post- adjustment stages. These constraints will specify which portions of the network should be traversed for some hop.

3 Methodology

In this paper, a graph representation is used consisting of a set of user nodes and a set of item nodes (with user nodes and item nodes not directly connected to each other). Each user node and item node has an associated activity, output and threshold. User and item nodes are connected via weighted edges where the weights on the edges represent ratings given to items by users (See Fig. 1). The activity of a user or item node a, for N nodes connected to the node a with non-zero weight, is calculated by:

$$activity_a = \sum_{i=1}^{N} x_i w_i$$

where x_i is the output of the node i that is connected to node a and w_i is the weight on the edge connecting node i to node a. The output of a user or item node is calculated based on that node's activity and a threshold value.

Spreading activation involves moving activation from one set of nodes to a second set of nodes. The terminology of a hop is used in this paper to define the activation spreading from one set of nodes to a second set of nodes. A hop involves the calculation of all node outputs in either the user set or item set, updating the associated activities and outputs of the nodes.

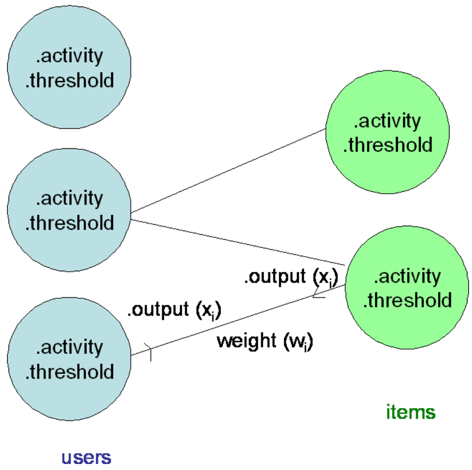

Fig. 1. Graph representation of Users, Items and Ratings

The initial hop performed is from the active user node to all item nodes connected to it (with non-zero weight). The activities and outputs of these item nodes are calculated and the nodes output a value. The second hop is from a set of item nodes to a set of user nodes, updating activities and outputs of user nodes. Two hops result in activating a set of user nodes constituting a user neighbourhood of the original active user node. The third hop, from user nodes to item nodes, provides item recommendations for the active user.

Much implicit information about users and groups can be extracted from the collaborative filtering data set and can be represented in this graph model. The initial experimental work in this paper only considers two of the simpler user and item features:

– *rated*: the number of items rated by some user in comparison to the maximum and minimum number of items rated by users.
– *avg-item-popularity*: the number of ratings received by some item in comparison to the maximum and minimum ratings received by items.

Based on the values of these features certain user and item nodes are identified prior to filtering and these nodes are used to constrain the spreading activation approach thus resulting in only some of the edges being considered and only portions of the graph being traversed.

4 Experiments

The experiments involve the comparison of two collaborative filtering approaches: a constrained spreading activation approach and a traditional memory-based collaborative filtering approach. A standard subset of the Movie Lens dataset is considered that contains the ratings of 943 users for 1682 movies. Weights on the network edges indicate the strength of like/dislike for an item where “dislike” can be viewed as an inhibitory or negative rating and “like” can be viewed as an excitatory or positive rating. Given that the original rating values in the Movie Lens dataset are all positive numbers, the approach adopted maps the ratings to positive and negative values to indicate positive and negative influences. The mapping chosen is to subtract 2.5 from all non-zero values, giving:

$\{0, 1, 2, 3, 4, 5\} \rightarrow \{0, -1.5, -0.5, 0.5, 1.5, 2.5\}$

A proportion of the data set is removed for testing and the metric of precision is used to compare the performance of the two approaches at different recall points. Precision is used because the spreading activation approach returns a ranking of recommended items, not prediction values that can be compared with actual values. In addition, work by Herlocker et al. [10] suggests that the level of accuracy provided by a metric such as Mean Absolute Error (MAE), which is often used in collaborative filtering evaluation, is unnecessary and that distinguishing between exact actual and predicted rating scores is not as important as measuring if the system correctly or incorrectly identifies “good” or “liked” items.

The traditional collaborative filtering approach uses Pearson Correlation to find correlated (similar) users. An adjustment is used in the Pearson Correlation calculation based on the number of items that users have rated in common (co-rated items) The motivation (as discussed in [13]) is that two users may receive a high correlation value but this might only be based on a small number of co-rated items. The adjustment ensures that users must have similar preferences over more than a few items to be considered highly correlated. The “nearest” neighbours of a user are selected using a low neighbour selection threshold, with any correlation value greater than 0.01 being considered. Although Breese [4] found that users with high correlation values (> 0.5) were more valuable in providing recommendations, work by Herlocker [9] using the Movie Lens data set found that for this data set such a high threshold sacrificed coverage and in addition, higher thresholds never improved the accuracy of predictions. They found that experiments with no threshold (using all correlation values > 0) always outperformed experiments with higher thresholds [9].

In the spreading activation approach to collaborative filtering, three stages corresponding to the three stages in the traditional memory-based collaborative filtering approach are used. Neighbours of some active user are found after activation has spread from the active user node to some set of item nodes (activating those items that the active user has rated) and from these item nodes to user nodes. At this stage the user nodes that have non-zero activity are the neighbours of the active user. When activation is spread again, from user nodes to item nodes, items not rated by the active user will be highlighted. These items are recommended to the user if the activity is sufficiently high. At any stage,

a node may be constrained in spreading its activation if it has been previously identified based on the analysis of user and item features.

5 Results

Figure 2 illustrates the precision recall graph for the two approaches and precision values are given in Table 1. Results were averaged over 30 runs. It can be seen that even with the limited sources of additional information included in this representation, the constrained spreading activation approach outperforms the traditional memory-based approach at all recall points other than at the first returned recommendation. These results were shown to be statistically significant using a 2-tailed paired t-test at p-value < 0.05.

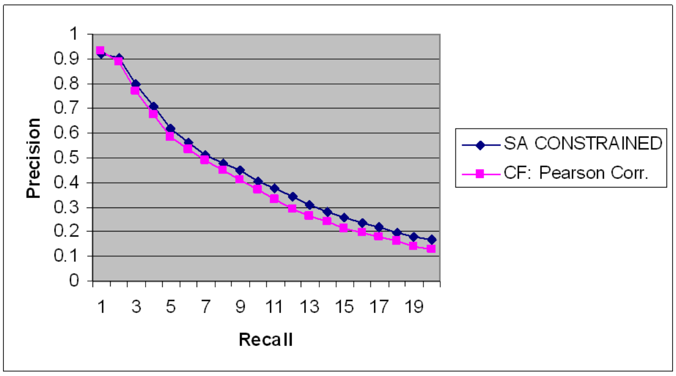

Fig. 2. Comparing Constrained Spreading Activation and Traditional Memory-based Approaches to Collaborative Filtering

On comparing a constrained and non-constrained version of the spreading activation approach (see Table 2), it was found that the constrained spreading activation approach outperformed the unconstrained approach at each recall point (the results were shown to be statistically significant using a 2-tailed paired t-test at p-value < 0.05.)

Recall Point	Constrained SA	CF: Pearson Corr.
1	0.9213	0.933
5	0.6194	0.585
10	0.4022	0.3688
15	0.2579	0.2136
20	0.1668	0.1273

Table 1. Comparing Constrained Spreading Activation and Traditional Memory-based Approaches to Collaborative Filtering

Recall Point	Constrained SA	Unconstrained SA
1	0.921277	0.909574
5	0.619362	0.611277
10	0.402163	0.397021
15	0.25792	0.251371
20	0.166755	0.16367

Table 2. Comparing Constrained and Unconstrained Spreading Activation Approaches to Collaborative Filtering

6 Conclusions

We have presented a spreading activation approach to collaborative filtering using features of the dataset to constrain the activation approach. Results show that even incorporating information on simple features shows improved performance over a traditional memory based approach. Future work will consider more complex user features, as well as group features, and will show that the graph representation and algorithm outlined and demonstrated in this paper can easily be extended to incorporate additional sources of evidence.

References

1. C.C. Aggarwal, J.L. Wolf, K.-L. Wu, and P.S. Yu. Horting hatches an egg: A new graph-theoretic approach to collaborative filtering. In *Proceedings of the Fifth ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'99), San Diego, CA*, pages 201–212, 1999.
2. M. Balabanovic and Y. Shoham. Fab: Content-based, collaborative recommendation. *Communications of the ACM*, 40(3):66–72, 1997.
3. C. Basu, H. Hirsh, and W. Cohen. Recommendation as classification: Using social and content-based information in recommendation. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*, pages 714–721, 1998.
4. J. Breese, D. Heckerman, and C. Kadie. Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann, July 1998.
5. P. Cohen and R. Kjeldsen. Information retrieval by constrained spreading activation on semantic networks. *Information Processing and Management*, 23(4):255–268, 1987.

6. F. Crestani and P.L. Lee. Searching the web by constrained spreading activation. *Information Processing and Management*, 36:585–605, 2000.
7. W. Croft, T. Lucia, J. Crigean, and P. Willet. Retrieving documents by plausible inference: An experimental study. *Information Processing and Management*, 25(6):599–614, 1989.
8. W.B. Croft. Combining approaches to information retrieval. In *Advances in Information Retrieval*, pages 1–36. Kluwer Academic Publishers, 2000.
9. J.L. Herlocker. *Understanding and Improving Automated Collaborative Filtering Systems*. Phd thesis, University of Minnesota, 2000.
10. J.L. Herlocker, J.A. Konstan, L.G. Terveen, and J.T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)*, 22:5 – 53, 2004.
11. Z. Huang, H. Chen, and D. Zeng. Applying associative retrieval techniques to alleviate the sparsity problem in collaborative filtering. *ACM Transactions on Information Systems*, 22(1):116–142, 2004.
12. Z. Huang, W. Chung, and H. Chen. A graph model for e-commerce recommender systems. *Journal of the American Society for Information Science and Technology*, 55(3):259–274, 2004.
13. M.R. McLaughlin and J.L. Herlocker. A collaborative filtering algorithm and evaluation metric that accurately model the user experience. In *Proceedings of the 27th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 329 – 336, 2004.
14. B. Mirza, B. Keller, and N. Ramakrishnan. Studying recommendation algorithms by graph analysis. *Journal of Intelligent Information Systems*, 20:131 – 160, March 2003.
15. C. O'Riordan and H. Sorensen. Multi-agent based collaborative filtering. In M. Klusch et al., editor, *Cooperative Information Agents 99, Lecture Notes in Artificial Intelligence*, 1999.
16. J. Palau, M. Montaner, and B. Lopez. Collaboration analysis in recommender systems using social networks. In *Cooperative Information Agents VIII: 8th International Workshop, CIA 2004*, pages 137 – 151, 2004.
17. A.M. Rashid, G. Karypis, and J. Riedl. Influence in ratings-based recommender systems: An algorithm-independent approach. In *SIAM International Conference on Data Mining*, 2005.
18. M.F. Schwartz and C.M. Wood. Discovering shared interests using graph analysis. *Communications of the ACM*, 36:78 – 89, August 1993.
19. U. Shardanand and P. Maes. Social information filtering: Algorithms for automating word of mouth. In *Proceedings of the Annual ACM SIGCHI on Human Factors in Computing Systems (CHI '95)*, pages 210–217, 1995.
20. W. Xi, E.A. Fox, W. Fan, B. Zhang, Z. Chen, J. Yan, and D. Zhuang. Simfusion: measuring similarity using unified relationship matrix. In *Proceedings of the 28th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 130 – 137, 2005.
21. G.-R. Xue, S. Huang, Y. Y., H.-J. Zeng, Z. Chen, and W.-Y. Ma. Optimizing web search using spreading activation on the clickthrough data. In *Proceedings of the 5th International Conference on Web Information Systems*, 2004.