

MULTI-INTEREST COMMUNITIES AND COMMUNITY-BASED RECOMMENDATION

Fang Wang

Pervasive ICT Research Centre, British Telecom Group, Orion 1/12, Adastral Park, Ipswich IP5 3RE, UK

Keywords: Community organisation, recommendation, association analysis.

Abstract: This paper introduces an approach to organising multi-interest communities in which a user may belong to more than one community. The interests of a user are first identified from the resources he handled and then refined through interest association analysis in order to remove false or redundant interests. To each identified interest topic, users who have this topic are clustered together, so a series of multi-interest communities are formed. Because members of a community may have interest in the topics of other communities, the formed communities are also connected with each other, resulting in a kind of community network indicating interest associations of groups of users. Based on formed multi-interest communities, users will receive useful recommendations within their own communities and from other related communities. This provides users opportunities to obtain information beyond their current interests so new interests of the users may be discovered. The multi-interest communities approach has been examined on the EachMovie data. The experimental results showed that the formed multi-interest communities were more cohesive and condensed when users were clustered according to their refined interest topics. The users also received much more recommendations based on multi-interest communities.

1 INTRODUCTION

The increasing use of the Internet brings out a new era of “online communities” or “virtual communities” that create a virtual space in which members do not necessarily meet and communicate with each other face-to-face. The formation of online communities has provided the opportunity for remote people to build social relationships and to exchange or share information with each other. Online communities have now been widely discussed in education, business and peer-to-peer computing. An important application of online communities is to provide users useful recommendations according to other like-minded users’ experiences. It is obvious that, in order to attain valuable recommendations, an essential factor is to find members with similar characteristics and construct proper online communities accordingly.

Many software platforms have perceived the importance and benefits of online communities and developed supportive services to encourage and facilitate group activities, e.g., synchronous or asynchronous communications between group members (Cassiopeia, Vignette, Webfair, e-groups).

Example application areas of these platforms include knowledge communities, business-to-business communities and customer-related communities. A few recent literatures in e-learning have investigated methods to create online learning communities because learning in collaboration, i.e., collaborative learning, is often more effective when students learn as a group (Seufert, et. al., 2002, Talavera and Gaudio, 2004). In addition to classifying students according to their subjects or demographic information derived from surveys, student clusters were learned probabilistically (e.g., via the Estimation Maximum algorithm) based on features mined from student interaction with the e-learning system. Within each formed community, useful recommendations were obtained by inspecting the reputation or content (e.g., meta-data) of each community member (Nijholt, 2002).

Interest-based communities have been intensively studied in peer-to-peer computing, as grouping remote peers according to their similarity will greatly improve the efficiency of resource search and sharing in a distributed environment. Generally, peers were given some attributes or objectives and joined into the same group when they

discovered other peers with the same defined attributes or objectives (Khambatti, et. al., 2004, Ogston, et. al., 2003). Peers were also linked together when they had similar access patterns to the same documents. This resulted in a kind of data-sharing graph which was small-world. Wang proposed a decentralised approach, self-organising communities, to organise distributed peers into a series of communities (Wang, 2002). Users with matching behaviour, e.g., one answered a query of another, were identified by a kind of middle agent and then grouped together. Without sophisticated user similarity calculation, self-organising communities successfully clustered users with similar interest or preferences. This work, however, assumed users had only one category of interest and belonged to one group each time. Self-organising communities have been introduced in an e-learning context to organise students into specific learning groups (Yang, et. al., 2004).

Most existing work of online communities attempt to allocate a user one most fitting community whose members have the maximum similarity. A user in such a community will receive pertinent information most interested by others. However, there is also a high probability that a user will never be able to discuss certain interesting information if this kind of information is widely unconcerned by other members of the community. Therefore, single-interest communities in which a user only belongs to one community tend to have a preference to 'popular' topics in which most community members are interested and exclude unpopular subjects that attract few members' attention.

This paper introduces a novel approach to organise multi-interest online communities, in which a user may belong to multiple communities. Multi-interest communities take users' manifold interests into full consideration and form communities based on user behaviour or activities, particularly the way they handle resources, e.g., access to resources or votes to resources. User interests are identified from the attributes of the resources they have handled, and further refined through interest association analysis. Users who have the same refined interest topic are grouped together, but have different association strengths to the group, reflecting their varied interest degrees. As a user may join into several communities, the formed communities develop into a community network at the same time. The directed connections between communities indicate how close one community is to another. As a consequence, recommendations can be made

within the communities and across communities that have close relationships.

The remainder of this paper is organised as follows. The next section introduces the formation of multi-interest communities. Recommendations based on multi-interest communities are explained in Section 3. Section 4 shows how the multi-interest communities approach worked on EachMovie data and how recommendations were accordingly made. The last section concludes this paper.

2 MULTI-INTEREST COMMUNITIES

The construction of multi-interest communities involves three key steps: user interest identification, community organisation and community network formation. Based on formed multi-interest communities, various kinds of recommendations can be made.

2.1 User Interest Identification

A user's interests are firstly identified from the resources he has handled and then refined via association analysis.

Identification of potential user interests

Resources or data in most practical applications such as documents or movies are usually associated with a set of well defined attributes or classes to summarise their general characteristics. A movie, for example, may have one or more genres including Action, Animation, Art_Foreign, Classic, Comedy, Drama, Family, Horror, Romance, and Thriller. When a user accesses a resource R_i , it usually indicates that this user is interested in the attributes, or at least some of the attributes of this resource. Because it is difficult to judge from this single access what exact attributes this user is interested in, the attributes of this resource only suggest a potential interesting topic to this user.

Suppose resource R_i has attributes or classes $\{a_1, a_2, \dots, a_n\}$. A user who has handled this resource may be interested in the combination of all these attributes, marked as T_i , or only part of the attributes – this needs to be further investigated as shown below. If all of the resources processed by a user are $\{R_1, R_2, \dots, R_m\}$, the potential interest topics of this user is the aggregation of those resources' attributes, noted as $\bigcup T_i, i=1, 2, \dots, m$.

Refining potential interests

In a potential interest set $\bigcup T_i$, a potential interest topic is obviously true if this topic has only one item (or attribute). A topic with more than one item may not be a valid interest to a user because it is possible that a resource R_i of this topic is accessed only because it involves a particular attribute that is of the user's real interest. However, if the user frequently accesses resources of the same topic, it is possible that this topic is a real one or at least near to a real interest of the user. When the potential interest topic is backed up by a good amount of resources of the same kind, this proves that this interest topic possesses enough Support and Confidence. Obviously the more resources a user handles, the more accurate would it be to verify the real interests of the user.

The Support and Confidence of a topic are defined similarly to those used in traditional association analysis (Han and Kamber, 2000). Suppose an interest topic T_i is the conjunction (combination) of attributes a_1, a_2, \dots, a_n . Support of T_i is then the percentage of resources processed by a user, which have topic T_i :

$$Support(T_i) = \frac{\text{Number of processed resources that have topic } T_i}{\text{Number of resources processed}} \quad (1)$$

Confidence of T_i in attribute a_j ($j=1,2,\dots, n$), is $c\%$ if $c\%$ of resources that has attribute a_j also has topic T_i . It is obvious that:

$$Confidence(T_i | a_j) = \frac{Support(T_i)}{Support(a_j)} \quad (2)$$

A potential interest topic T_i is said to be a valid interest of a user if the confidences on all of the attributes a_j satisfy a confidence threshold γ .

If T_i is proven to be invalid from its confidence calculation, the combinations of part of the attributes a_j may still be valid interests if they have enough confidences on their attributes. A potential interest topic will be examined in this way until a valid topic is obtained or there is only one item/attribute left. The final interests of a user are composed of all of the valid interest topics after examination on $\bigcup T_i$, including interests with either a single attribute or a combination of multiple attributes.

2.2 Community Organisation

The organisation of communities is relatively straightforward after obtaining refined interests of all of the users: to each refined interest topic T_i , a

community of this topic will be created to include all users who possess this topic. By grouping users that have the same valid interest topic together, a series of communities are obtained. Because a user usually has more than one interest topic after interest refining, he will join in several communities at the same time. Every user has an association strength associated with each of his communities. The association strength indicates the interest degree of a user to an interest topic or a community. In real applications, the strength may be defined as the number of resources handled by this user within a community or the confidence level of the user in a community.

2.3 Community Network Formation

A community network illustrates the relationships between communities. To be exact, it reflects the correlations of the interests of groups of users. If most members in a community A, for instance, are also members of a community B, this suggests that most members interested in the topic of community A are also interested in the topic of community B. As a result there is a strong connection from community A to B and the connection strength indicates the closeness of A to B.

However, even if most members of A belong to B, it is possible that these members are only a small portion of the members of B. So the connection from B to A may be much weaker than that from A to B. The resulting community network is hence a directed and asymmetric graph.

By observing community correlations, we will not only get knowledge of how individual users share common interests together, but also how they share the same resource access pattern and how different kinds of interest topics are indirectly related because of their groups of users. The latter will provide valuable information for recommendation across related communities.

3 RECOMMENDATION BASED ON MULTI-INTEREST COMMUNITIES

The formation of multi-interest communities can be useful to a series of practical applications. For instance, by clustering users into groups, it would be much easier for them to share information and build social relationships together. Personalised services can also be provided to individual users or groups of

users according to their identified interests. Another typical application of multi-interest communities is recommendation of valuable information to a user or groups of users according to other similar or related users' experiences. Two kinds of recommendations can be made based on multi-interest communities: intra-community recommendation and inter-community recommendation.

Intra-community recommendation

Intra-community recommendation takes place within each community. As all of the members in a community are interested in the same topic, it is possible that a member will be interested in the resources of the same kind and most accessed by other members. In contrast to most collaborative filtering techniques that do not discriminate recommended data about their relative popularity, this paper ranks resources of a community for recommendation. This is similar to the work reported in (Lawrence, et. al., 2001). However, in addition to resource popularity, the recommendation provided in this paper is also dependent on how much information a user to be recommended requires and how close the user is associated with his community (i.e., how much the user is interested in the community topic). Therefore, the recommendation $\mathfrak{Rec}(r_j, u_i)$ of a resource r_j in community C_k to a user is a function of the user's requirement, the user's relationship to the community and the popularity or usefulness of the resource deemed by the other users of the community. This is illustrated by the following equation:

$$\mathfrak{Rec}(r_j, u_i) = \text{Req}(u_i) \cdot \omega(u_i, C_k) \cdot \text{Pop}(r_j, C_k) \quad (3)$$

where $\text{Req}(u_i) \in [0,1]$ indicates the recommendation requirement of user u_i . It has a maximum value 1, which means that the user welcomes all recommendations. $\text{Req}(u_i) = 0$ suggests that the user does not want any recommendations. $\omega(u_i, C_k)$ is the association strength of a user to his community C_k . It can be decided explicitly by the user or implicitly from the user's activities in the community, e.g., how many resources the user has accessed. $\text{Pop}(r_j, C_k)$ is the popularity or usefulness of a resource in community C_k . It can be measured from the resource's access frequency or voted scores of the community members. Resource r_j will be recommended to user u_i if the final $\mathfrak{Rec}(r_j, u_i)$ obtained from equation (1) is greater than a threshold ϵ_1 .

Inter-community recommendation

In addition to recommendation among community members, groups of users are able to receive useful information recommended from other related communities. Due to the introduction of inter-community recommendation, users will not only receive useful information from similar-minded peers, but also information of different kinds but most referred by other related peers. Inter-community recommendation provides a user an opportunity to obtain information beyond his current communities and as a result, new interests of this user may be identified.

In inter-community recommendation, a community A will receive recommendations from community B only if there is a connection from A to B, suggesting that some users interested in A's topic have also interest in the topic of B. Again the resources of a community are ranked according to their popularity. The actual recommendation $\mathfrak{Rec}(r_j, C_k, C_l)$ of a resource r_j in community C_l to community C_k depends on the popularity of resource r_j in C_l , the association strengths or closeness from C_k to C_l and the member requirements of community C_k , as illustrated below:

$$\mathfrak{Rec}(r_j, C_k, C_l) = \text{Req}(C_k) \cdot \omega(C_k, C_l) \cdot \text{Pop}(r_j, C_l) \quad (4)$$

where $\text{Req}(C_k) \in [0,1]$ indicates the requirement of community C_k . $\omega(C_k, C_l)$ is the association strength of the connection from community C_k to community C_l and $\text{Pop}(r_j, C_l)$ is the popularity of resource r_j in community C_l . Resource r_j will be recommended to community C_k if $\mathfrak{Rec}(r_j, C_k, C_l)$ is greater than a threshold ϵ_2 .

When community C_k receives a recommended resource r_j , it will not disseminate this resource to every user in the community. Resource r_j will only be forwarded to users that may have interest in it. This will be verified by equation (3) as introduced above, that is, recommendation of r_j to a particular user u_i in community C_k is also decided by the user's individual requirement and his closeness to the community C_k .

4 EXPERIMENTAL SIMULATION

Multi-interest communities have been tested on the EachMovie data provided by the DEC systems

research centre (EachMovie). EachMovie data were collected from 72916 users who entered a total of 2811983 numeric ratings for 1628 different movies (films and videos). Every movie falls into one or more of 10 genres: Action, Animation, Art_Foreign, Classic, Comedy, Drama, Family, Horror, Romance and Thriller. User votes were recorded in the following format:

Person_ID	Movie_ID	Score	Weight	Modified: Date/Time
-----------	----------	-------	--------	---------------------

Here, Score and Weight are numerical numbers between 0 and 1, i.e., $0 < \text{Score}, \text{Weight} \leq 1$. In particular, the Score maps linearly to the zero-to-five star rating, that is, it has a value of 0, 0.2, 0.4, 0.6, 0.8 or 1. Weight indicates whether the person rated a movie as zero to five stars (weight = 1) or "sounds awful" (weight < 1). (Most "sounds awful" weights are 0.2, but for historical reasons about 10% are 0.5.)

The EachMovie data set was initially built for examining collaborative filtering techniques. It was used in this paper to particularly test the formation of multi-interest communities and the recommendation made on multi-interest communities.

User Interest Identification

5000 random users were chosen from the EachMovie data for the examination presented in this paper. The interests of these users were first identified from their voted movies as introduced in Section 2.1.

Here we take user 10 as an example. This user voted in total 88 movies, belonging to 27 kinds of genres, as shown in Figure 1. This suggests that the potential interests of user 10 included those 27 genres. For 5000 users, in total 76 kinds of potential interest topics were identified from their votes.

In the potential interest topics of a user, some of them may be false or subject to more general topics so they should be removed from the user's real interests. In order to identify those invalid interest topics, refining rules introduced in Section 2.2 were employed to calculate the Support and Confidence values for all potential interest topics. Figure 2 shows the Support values of user 10 on ten movie genres. Figure 3 shows the confidences of user 10's 27 potential interest topics on movie genres. In this experiment, the confidence threshold γ was set as 0.1. Interest topics with a genre confidence lower than 0.1 suggest that these topics are not of major interest to the user. After interest refining, 7 valid interests were obtained from user 10's 27 potential

interests. The final valid interest topics include (Comedy Romance), (Drama), (Art_Foreign), (Comedy), (Animation Family), (Action Thriller) and (Action), as indicated in Figure 1.

It is worth noting that (Action Thriller) is treated as a separate valid interest though it is logically subject to (Action). It is because this category received many votes so the confidences on both genres of Action and Thriller are high. This suggests that user 10 has strong interest in this specific category, though he has also a broad interest in general Action movies. Another category (Animation Family) only received 2 votes from user 10. However, in the EachMovie data, there were in total only two movies of this category and user 10 voted for both of them. This category represents user 10's particular interest in Animation and Family and is hence recognised as a positive interest of this user.

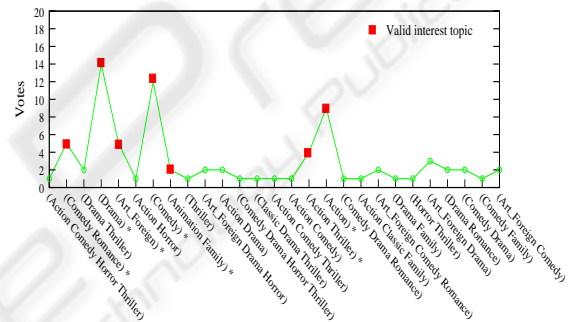


Figure 1: Votes of User 10.




Figure 2: User 10's Support on movie genres.

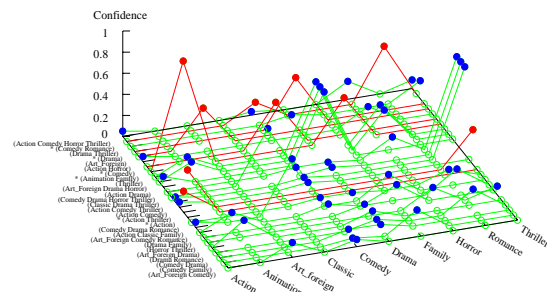


Figure 3: Confidences of user 10's potential interest topics on movie genres.

Through the interest refining process, the total number of interest topics of 5000 users was reduced from 76 to 58. On average a user has 9 interests. Some users showed broad interests with as many as 21 interest topics, whereas some others had only 1 interest topic. Figure 4 shows the distribution of the number of user interests. This result suggests that the number of people that own k different interests does not vary when the number of interests is smaller than 9. That is, the probability that someone has $k < 9$ interest topics is independent of k . Beyond this point, for intermediate values of k ($9 < k < 16$), it is likely that someone chosen at random belongs to exactly k communities. Finally, as k goes beyond 16, the probability that a user belongs to $k > 16$ communities decays exponentially fast in k .




Figure 4: Distribution of number of user interests.

Formation of multi-interest communities

For each refined interest topic, all of the users that showed this interest were suggested to join into the same group of this topic. This finally resulted in 58 communities. The number of members of these communities ranged from 1 to 8449. The top ten communities that have the most members are (Drama), (Action), (Comedy), (Action, Thrill), (Animation Family), (Comedy Romance), (Drama Thriller), (Action Drama), (Thriller) and (Horror Thriller). Actually these communities were also the most popular groups before user interest refining but had more members.

For a comparison of the communities made before and after interest refining, the community cohesiveness and the average user strength of each community were calculated for both cases. Here the weighted similarity (Steinbach, et. al., 2002), which is the squared length of the community centroid, was used to test the internal cluster similarity, as shown in equation (5).

$$\text{Centroid}(C) = \frac{1}{|C|^2} \sum_{u_1, u_2 \in C} \text{cosine}(\dot{u}_1, \dot{u}_2) \quad (5)$$

where users u_1 and u_2 are members of community C , \dot{u}_1 and \dot{u}_2 are voting vectors of users u_1 and

u_2 respectively. User strength was measured as the number of votes voted for the movies in each community, normalised by the total number of movies in the community.

Table 1 lists the average community cohesiveness and average user strength in communities before and after user interest refining. It proves that the community cohesiveness, or intra-community user similarity, was improved from 4.11 to 7.32 when users are clustered according to their refined interests. At the same time, the average user strength which indicates user closeness to their communities was increased from 0.35 to 0.45. These results suggest that more condensed communities were obtained when interest refining was employed to purify user interest topics. By removing relatively irrelevant community members and loose communities, the resulting communities had improved closeness among community members and between community and their members.

Table 1: Community cohesiveness and average user strength before and after interest refining.

	Community cohesiveness	Average user strength
Before interest refining	4.11	0.35
After interest refining	7.32	0.45

Community Network

Depending on the relationships of community topics, 58 communities formed a complex network. For any two communities that have connections, which means those communities have common users, the strength of connection from community C_k to C_l was calculated as follows:

$$\omega(C_k, C_l) = \frac{\sum_{u \in C_k \cap C_l} \omega(u, C_k)}{\sum_{U \in C_k} \omega(U, C_k)} \quad (6)$$

where the denominator is the sum of the strengths of all users in community C_k and the numerator is the strength sum of the common users of community C_k and C_l . This community connection strength shows how much the users in community C_k are interested in the topic of community C_l . As shown by equation (6), the connections between C_k and C_l are directional, depending on the percentage of the common users in each community.

The community network of EachMovie communities is illustrated as a directed graph as shown in Figure 5. Only 2600 relationships were presented among 3306 possible connections between

58 communities. Some communities had no direct connection to each other at all, such as communities (Animation Comedy) and (Classic), and (Romance) and (Drama Horror), which means those communities had no common users. Some communities showed a kind of hierarchy, such as the connection from (Animation Comedy) to (Comedy), as all of the users of the former were also members of the latter. Hierarchical connections usually have high connection strengths.

Close relationships have been found between popular communities, such as connections from (Art_Foreign Classic) to (Action) and from (Action Horror Thriller) to (Drama), though those communities are semantically non-related. Actually the most popular 10 communities were nearly fully connected with each other and these communities shared at least half of their members. The connections between popular communities are also asymmetric. For example, the relationships from (Action) to (Art_Foreign Classic) and from (Drama) to (Action Horror Thriller) were much weaker than those of the other direction.




Figure 5: Community network.

Community-based Recommendation

Intra-community Recommendation

Intra-community and inter-community recommendations can now be made based on the organised communities and community network. User 10 was again taken as an example to show how those two kinds of recommendations could be made for him. It is assumed that this user welcomed all recommendations, that is, $Req(u_{10}) = 1$. As introduced in Section 4.1, user 10 belonged to seven communities, including (Comedy Romance), (Drama), (Art_Foreign), (Comedy), (Animation Family), (Action Thriller) and (Action). The top 10 movies in each community were selected for recommendation, as listed in Appendix A, except those movies marked with *, which were already viewed and voted by this user. As a result, there

were in total 97 movies sent to user 10 for consideration.

Inter-community recommendation

The seven communities where user 10 belonged to connected many other communities. This was because the other members of those seven communities showed strong or weak interests in other movie genres. Some communities were linked by more than one of those seven communities in the formed community network with varied connection strengths. Figure 6 shows the average connection strengths of the connected communities. As shown in Figure 6, there were around 10 other movie communities most linked by user 10's communities. Those ten communities were hence selected for inter-community recommendation to user 10. The top 10 movies in these strongly related communities are listed in Appendix B. Most of the movies were not viewed or voted by user 10. Actually, among those selected 10 other communities, 4 of them were new to user 10 as this user had no votes to the movies of these communities. User 10 had only one or two votes to the other six connected communities.

Depending on how many recommendations user 10 would accept, we could further inspect this user's real interests or even find out his new interests that had not been shown in this user's votes before. This work, however, is impossible to do due to the anonymity of EachMovie users. We hope to do more tests on real users to complete this work in our near future.




Figure 6: Average connection strengths of related communities of user 10.

5 CONCLUSIONS AND FUTURE WORK

This paper introduces an approach to organising multi-interest communities in which a user may belong to more than one community. The interests of a user are first identified from the resources he

handled and then refined through interest association analysis in order to remove false or redundant interests. To each identified interest topic, users who have this topic are clustered together, so a series of multi-interest communities are obtained. Because members of a community may have interest in the topics of other communities, the formed communities are also connected with each other, resulting in a kind of community network indicating interest associations of groups of users. Experiments on the collected EachMovie data showed that the formed multi-interest communities were more cohesive and condensed when users were clustered based on their refined interest topics.

Intra-community and inter-community recommendations can be made based on formed multi-interest communities. The former recommends to a user popular resources deemed by other community members within a community, whereas the latter suggests resources of other categories but most interesting to the other members of a community. Consequently a user will receive information within and beyond his identified interests. From his responses to the recommendations, e.g., accept or reject some of them, the user's real interests will be further identified and even new interests of the user may be discovered.

The recommendation approach presented in this paper is more suitable to users that welcome all information recommendations. Some users in reality may only wish to receive carefully selected information, which means their recommendation requirements $Req(u) < 1$. How to tailor recommendations according a user's requirements and preferences will be studied as our future work. In addition, a user's interest to a certain topic or resource will be divided into positive and negative, instead of all positive interests as shown in this paper. Because the users of the EachMovie data are anonymous, it is difficult to judge whether they would like the groups allocated to them and the recommendations suggested by our approach. The proposed multi-interest communities and community-based recommendation will be further tested on real users so that the users' feedback will be used to examine and improve the approach presented in this paper.

ACKNOWLEDGEMENTS

The author is grateful to the support from BT Long Term Research Venturing, colleagues in Nanjing University of China and the National Natural Science Foundation of China Grant No.60402027.

REFERENCES

- Cassiopeia <http://www.cassiopeia.com/>
 EachMovie <http://research.compaq.com/SRC/eachmovie/>
 e-groups <http://www.egroups.com/>
 Han, J. and Kamber, M., 2000. Data Mining: Concepts and Techniques, Morgan Kaufmann Publishers.
 Iamnitchi, A., Ripeanu, M. and Foster, I., 2004. Small-World File-Sharing Communities, The 23rd Conference of the IEEE Communications Society. Hong Kong.
 Khambatti, M., Ryu, K.D. and Dasgupta, P., 2003. Structuring Peer-to-Peer Networks using Interest-Based Communities. *1st International Workshop on Databases, Information Systems, and Peer-to-Peer Computing*. Berlin, Germany, pp. 48-63.
 Lawrence, R.D., Almasi, G.S., Kotlyar, V., et. al. 2001. Personalization of Supermarket Product Recommendations. *Data Mining and Knowledge Discovery* 5(1-2):11-32.
 Nijholt, A., 2002. Computer-facilitated community building for E-learning, *Proc. IEEE Inter. Conf. on Advanced Learning Technologies*. Kazan, Russia, pp.541-543.
 Ogston, E., Overeinder, B., van Steen, M., and Brazier, B., 2003. Group Formation Among Peer-to-Peer Agents: Learning Group Characteristics. *Inter. Workshop on Agents and Peer-to-Peer Computing*, pp. 59-70.
 Seufert, S., Lechner, U. and Stanoevska, K., 2002. A reference model for online learning communities. *Inter. J. on E-Learning*, Jan-Mar, pp.43-55.
 Steinbach, M., Karypis, G. and Kumar, V., 2000. A comparison of document clustering techniques, *KDD Workshop on Text Mining*.
 Talavera, L. and Gaudioso, E., 2004, Mining student data to characterize similar behaviour groups in unstructured collaboration spaces. *Workshop on Artificial Intelligence in CSCL, 16th European Conference on Artificial Intelligence*. pp.17-23.
 Vignette <http://vignette.com/>
 Webfair <http://www.webfair.com/>
 Wang, F., 2002. Self-organising communities formed by middle agents. *Proc. of the 1st Inter. Conference on Autonomous Agents and Multi-Agent Systems*, Bologna, Italy, pp 1333-1339.
 Yang, F., Shen, R. and Han, P., 2004. A novel self-organizing e-learner community model with award and exchange mechanisms, *Journal of Zhejiang University Science*, 5(11): 1343-1351.

Appendix A: Intra-community recommendations for user 10.

Movies marked with * were already viewed and voted by this user.

	(Comedy Romance)	(Drama)	(Art_Foreign)	(Comedy)	(Animation Family)	(Action Thriller)	(Action)
1	Late Bloomers	Hard Eight	Identification of a Woman	The Full Monty	* Toy Story	Aliens (1986)	Air Force One
2	Love and Other Catastrophes	Rosewood	The Eighth Day	Raising Arizona (1987)	The Lion King	In the Line of Fire	The Terminator (1984)
3	Groundhog Day	In the Company of Men	Jean de Florette (1986)	A Fish Called Wanda (1988)	Winnie the Pooh and the Blustery Day	Breakdown	Die Hard (1988)
4	When Harry Met Sally... (1989)	Schindler's List	Manon of the Spring (1986)	This Is Spinal Tap (1984)	Snow White and the Seven Dwarfs	Operation Condor (Feiyang Gaiwak)	Terminator 2: Judgment Day
5	Strictly Ballroom	The Shawshank Redemption	The City of Lost Children	Local Hero (1983)	* Beauty and the Beast	Heat	* The Rock
6	Sleepless in Seattle	Sling Blade	Le Colonel Chabert	Monty Python's Life of Brian	The Fox and the Hound (1981)	Star Trek II: The Wrath of Khan (1982)	Speed
7	My Best Friend's Wedding	* Lone Star	Paris Was a Woman	Back to the Future (1985)	Robin Hood (1984)	Clear and Present Danger	Blade Runner (1982)
8	The American President	Ran (1985)	Madame Butterfly	My Favorite Year (1982)	Tim Burton's The Nightmare Before Christmas	Face/Off	Hoodlum
9	* The Truth about Cats and Dogs	Hamlet (1996)	Le Confessionnal	Living in Oblivion	The Hunchback of Notre Dame	* Independence Day (ID4)	Full Metal Jacket (1987)
10	* While You Were Sleeping	Traveller	Delicatessen	Grosse Pointe Blank	Casper	Highlander (1986)	The Big Blue (1988)

Appendix B: Intra-community recommendations for user 10.

Movies marked with * were already viewed and voted by this user.

	(Thriller)	(Horror)	(Horror Thrill)	(Family)	(Drama Thriller)	(Comedy Family)	(Action Drama)	(Romance)	(Animation)	(Drama Romance)
1	The Game (1997)	Paradise Lost: The Child Murders at Robin Hood Hills (1996)	* The Silence of the Lambs	Daniel Defoe's Robinson Crusoe	Mother Night	Babe	Contact	William Shakespeare's Romeo and Juliet (1996)	The Wrong Trousers	Love Jones
2	Nightwatch	The Shining (1980)	Jaws (1975)	Mary Poppins (1964)	* The Usual Suspects	Liar, Liar	Miller's Crossing	Before Sunrise	A Close Shave	The Whole Wide World
3	Reservoir Dogs	An American Werewolf in London (1981)	Scream	Pinocchio (1940)	Once Upon a Time in America (1984)	Mrs. Doubtfire	Braveheart	One Fine Day	A Grand Day Out	Somewhere in Time (1980)
4	Bound	Freeway	The Abyss (1989)	Willy Wonka and the Chocolate Factory (1971)	Murder in the First	Home Alone	GoodFellas	Benny & Joon	Wallace & Gromit: The Best of Aardman Animations	Infinity
5	Unforgiven	Bram Stoker's Dracula	Copycat	Fly Away Home	The Frighteners	The Santa Clause	Fire Down Below	Pretty Woman	Ghost in the Shell (Kokaku Kidotai)	Blue Sky
6	Crimson Tide	Interview with the Vampire	Cape Fear (1991)	A Little Princess	Sleepers	Matilda	G.I. Jane	It Could Happen to You	Heavy Metal (1981)	Chasing Amy
7	Ransom	The Howling (1981)	Mary Shelley's Frankenstein	The Jungle Book	* Tombstone	Muppet Treasure Island	The Long Kiss Goodnight	For the Moment	The Transformers: The Movie (1986)	Jerry Maguire
8	Red Rock West	The Prophecy (God's Army)	Mute Witness	Jumanji	The Client	The Stupids	* Seven	Bed of Roses	n/a	Some Kind of Wonderful (1987)
9	The Grifters	A Nightmare on Elm Street (1984)	Body Snatchers	Cool Runnings	Outbreak	The Big Green	Donnie Brasco	Pie in the Sky	n/a	Ghost
10	Dolores Claiborne	Cat People (1982)	In the Mouth of Madness	Andre	The Firm	Jack	True Romance	Dirty Dancing	n/a	Up Close and Personal