Exploiting Social Ties for Search and Recommendation in Online Social Networks – Challenges and Chances

Kerstin Bischoff
L3S Research Center / Universität Hannover
Appelstrasse 4
30167 Hannover, Germany
bischoff@L3S.de

ABSTRACT

Online social networking is a huge trend. On *Facebook* or *MySpace* people connect with their friends or make new friends. They form new (indirect) connections by reading and adopting from other peoples' blogs or tweets. Similarly, in tagging systems like *Delicious* or *Flickr* people share tagged resources with friends or unknown, similar users. Offline social networks have long been studied in sociology, epidemiology, etc. However, the new online networks offer new ways to revisit old theories as well as to find emerging trends with respect to information diffusion and sharing in the Web 2.0. The goal is to explore and exploit relational ties in a way to enable the mining of useful knowledge and effective information propagation/diffusion. Assuming no or only potential ties explicitly given, the focus is first on the analysis of collaborative tagging and its potentials for user profiling, recommendations and search. Some first related studies, approaches and ideas for future work address the identification and exploitation of weak and strong ties in online networks.

Categories and Subject Descriptors

H.3.1 [**Information Storage and Retrieval**]: Content Analysis and Indexing; H.3.4 [**Information Storage and Retrieval**]: Systems and Software; H.3.5 [**Information Storage and Retrieval**]: On-line Information Services

General Terms

Experimentation, Human Factors, Algorithms

Keywords

social network analysis, social ties, collaborative tagging, search, recommendation

1. INTRODUCTION

With the advent of the Web 2.0 online social networking has become a huge trend. On platforms like *Facebook* or *MySpace* people connect with their friends or make new friends. They form new (indirect) connections by reading and adopting from other peoples' blogs or tweets. Similarly, in collaborative tagging systems like *Last.fm, Delicious* or *Flickr* people share bookmarks and tagged resources with friends or unknown, similar users. Reasons are manifold: staying in touch, socializing, finding answers/experts, share resources and knowledge, etc. Offline social networks have long been studied under the perspective of sociology, epidemiology, and even thermodynamics. However, online networks offer new ways to revisit old theories as well as to find emerging trends in information diffusion and sharing in the Web. For search and recommendation, the topic of tie strength is interesting. In 'real' social networks, strong (i.e. family, close friends) and weak ties (loose acquaintances) have been found to show different characteristics e.g. with respect to services offered (see section 3). As McAfee [19] pointed out, different kinds of ties, if supported by the right technology, may offer different potential benefits for information exchange and collaboration (adapted from [19]):

- Strong ties: Collaboration in a closed group, e.g. co-workers (BSCW, CVS, Wikis)
- Weak ties: Innovation, Non-redundant information, Network bridging (Email, Social networking systems)
- Potential ties: Efficient search, Tie formation (Blogosphere, Bulletin boards, Folksonomies)
- No/Absent ties: Collective Intelligence, statistical patterns (Folksonomies, Prediction Market, Question Answering)

The goal is to explore and exploit relational ties in a way to enable the mining of useful knowledge and effective information propagation/diffusion so that people are provided the information they need. Focusing on absent, potential and weak ties, we will present first research results, algorithms and ideas for future work. To better understand 'collective intelligence' expressed via tags, i.e. statistical patterns found in folksonomies, we start by describing an analysis of different tagging systems and summarize some experiments on how to exploit social annotations to enrich metadata for multimedia resources. Web 2.0 tools and environments like the personalized Internet radio and social music network *Last.fm*[1] have made collaborative tagging so popular: any user can assign freely selected words, in the form of keywords or category labels, to shared content – thereby describing and organizing these resources. As a result, a

GvD Workshop'10, 25.-28.05.2010, Bad Helmstedt, Germany.

[1]http://last.fm

huge amount of manually created metadata describing all kinds of resources as well as user interests is now available. Such semantically rich user generated annotations are especially valuable for recommending, searching and browsing multimedia resources, such as music, where these metadata enable retrieval on the newly available textual descriptions represented by tags. However users' motivations for tagging resources, as well as the types of tags differ across systems. These tags represent quite a few different aspects of the resources they describe [2] and more research is needed on how these tags or subsets of them can be used effectively for user profiling and search in social networks. Weak and strong ties and the corresponding topics of information diffusion and social search are considered next. Weak ties are often 'bridges' connecting different communities, thus bringing new information (e.g. job seeking). Strong ties offer mutual support, trust, but likely share knowledge, preferences, values and friends. The section will cover the related issues of identifying tie strength or characterizing friendship relations in online networks. It introduces recent ideas to analyze and exploit such networks to improve common approaches to search and recommendation.

2. TAGS FOR USER PROFILING, SEARCH AND RECOMMENDATIONS

As they offer a promising way to estimate similarity between resources, users and resources or between different users, the usefulness and reliability of tags is important for many search and recommendation algorithms. In the first section the focus will be on analyzing tag usage patterns and their implications for user profiling, search and recommendation. Then, we will present approaches exploiting tags to enrich resources or user profiles with additional information – music moods and themes as well as picture moods.

2.1 Usefulness of tags for profiling and search

Here we present results of an in-depth study [2] of tagging for different kinds of resources and systems – Web pages (*Delicious*), music (*Last.fm*) and images (*Flickr*). Tags serve various functions based on system features like resource type, tagging rights, etc. [17], and not all these tags are equally useful for user profiling or for interpersonal retrieval. For being able to improve tag based user profiling and search, we first need to know how tags are used and which types of annotations we can expect to find along with resources. [17] identifies organizational motivations for tagging, other more social motivations include opinion expression, attraction of attention, and self-presentation.

Our analysis revealed the necessity and usefulness of a common tag classification scheme for different collections, which allows the comparison of the types of tags used in different tagging environments. For example, the distributions of tag types strongly depend on the resources they annotate: for *Flickr* and *Delicious* the most frequently used tags (50% of the cases) refer to topic concepts (i.e. what the resource is about), while for *Last.fm*, type-related tags (e.g. genres) are the most prominent ones. Other interesting results of this analysis refer to the added value of tags to existing content: More than 50% of existing tags bring new information to the resources they annotate. For the music domain this is even the case for 98.5% of the tags. Especially for multimedia data, such as music, pictures or videos, the gain provided by the newly available textual information is substantial, since with most prominent search engines on the Web, users are currently still constrained to search for multimedia using textual queries. A large amount of tags is also accurate and reliable; in the music domain for example 73.01% of the tags also occur in online music reviews written by experts.

Regarding search, our studies showed that most of the tags can be used for search and that in most cases tagging behavior exhibits approximately the same characteristics as searching behavior. However, some noteworthy differences have also been observed. Namely, for the music domain, the usage context (i.e. situation suitable for listening to a particular song – e.g. "pool party") is very useful for search, yet underrepresented in the tagging material. Similar, for pictures and music opinions or qualities (e.g. characteristics, moods) queries occur quite often, although people tend to neglect this category for tagging. Clearly, supporting and motivating tags within these categories could provide additional information valuable for search.

2.2 Knowledge mining from tags

Our analysis [2] showed some clear gap between the tagging and the querying vocabulary for music as well as pictures. For pictures, a large portion of tags refer to location information. However, queries targeting images much more often name subjective aspects, e.g. "scary", "rage" or "funny". For music, tags predominantly name the genre (i.e. type), though when searching for music, the majority of queries falls into these categories: 30% of the queries are theme-related, 15% target mood information. In this section we will shortly present an approach for detecting moods and themes for songs [2, 3, 5] and emotions in photos [4]. It relies on collaborative tagging and aims at bridging exactly this gap identified. The methods proposed can be used in various ways: as part of an application where the recommendations are presented to the user for selection, for indexing and thus enriching the metadata indexes to improve searchability, or to automatically create mood-based picture catalogs.

2.2.1 Datasets

For our experiments we gathered data from several sources (for details please refer to [3, 4]):

- From the *Allmusic.com* website we collected the labels of 178 different moods and 73 themes together with music tracks that fall into these categories.
- For 13,948 songs obtained from *Allmusic.com*, we could get user tags from the social music platform *Last.fm*.
- To investigate whether lyrics can provide added value in the task of mood and theme recommendation, we also obtained the lyrics for our tracks from *lyricsdownload.com* and *lyricsmode.com*.
- For the purpose of deriving mood labels for pictures, we manually selected *Flickr* groups that correspond to the emotion/mood labels in the hierarchy of human emotions presented in [25]. The taxonomy comprises the six primary emotions "Love", "Joy", "Surprise", "Anger", "Sadness", and "Fear", each of which has more fine-grained secondary emotions. We found

corresponding *Flickr* groups for 17 out of the 25 secondary emotions, including the six primary emotion labels. For all pictures in the identified groups we downloaded all associated user tags from *Flickr*.

2.2.2 Picture Mood Recommendation Algorithm

The assumption on which we base our recommendation is that the existing tags attached to music songs or photos can possibly provide information regarding the corresponding mood (or theme). The user given tags are used as input features for training a classifier over all mood or theme classes. Here we also make use of the Weka[2] implementation for the Naïve Bayes Multinomial classifier, which produces for all pictures in the test set probability distributions over all classes of moods. For music, training and evaluation is done using the *Allmusic.com* ground truth, created manually by music experts. All distinct tags span the feature vector of a song, weighted by the frequency with which the tag occurred for a song. Similarly, all pictures pertaining to a specific mood class (i.e. social group on *Flickr*) represent the positive training examples, while pictures taken randomly from the rest of the classes build up the set of negative examples. The number of positive and negative examples for a class was equally balanced. In order to ensure a fair classification of the data, here all tags related to a mood or emotion were deleted using WordNet synset information.

2.2.3 Experiments and Results

Music. Different experimental runs were performed using either tags, lyrics, or both plus varying classes to be predicted. Since the 178 mood terms from *Allmusic.com* are hardly distinguishable for a non-expert, these labels were manually mapped to the hierarchy of basic human primary and secondary emotions (see 2.2.1). This resulted in 22 secondary mood terms and six primary classes. For themes, the eleven classes were reduced to nine by applying a WordNet based clustering accounting for word overlap in synsets. The best performing methods are those using tags as input features, while classification based on lyrics performs worst. Combining tags and lyrics achieves good results, sometimes even slightly outperforms tag-based method. While incorporating lyrics features led to good results for genre [3], they do not seem to be indicative of the mood of a song. For themes, there is a slight, yet rarely significant, effect. For the case of theme recommendations, given the original eleven themes we achieve a $H@3$ of 0.80 based on tags and lyrics. The best results, $H@3$ of 0.88, are achieved for the algorithm using a combination of tags and lyrics as features and applying a WordNet synonymy based clustering on the theme classes. Compared to themes, mood recommendations do not perform as well when using many classes, achieving only a $H@3$ of 0.64 for the manually clustered 22 secondary emotions. Reducing the number of clusters to the 6 first level classes of basic human emotions boosts the performance considerably and for the best method using tags and lyrics as input features we achieve a $H@3$ value of 0.89. Micro-evaluating results per specific class, shows that some classes are relatively easy to recommend. Others may require special attention or some level of disambiguation. In general, those class labels that are harder to recommend appear more ambiguous with the corresponding annotations being mostly subjective. Themes like "Late Night" or "Summertime" strongly depend on each person's individual associations.

[2]http://www.cs.waikato.ac.nz/~ml/weka/

Pictures. A first set of experiments aimed at recommending mood labels corresponding to the primary human emotions. In this case, the classes to be learned by the classifiers consisted of the union of all data belonging to all underlying secondary emotions (e.g. the Love class comprises all data gathered from the *Flickr* groups for Affection, Lust and Longing). For the experimental runs on secondary emotion label recommendations, each secondary emotion class represented a class to be learned. We perform a 10-fold cross validation and evaluate the performance of our method according to standard IR metrics. All recommendations corresponding to the primary human emotions achieved very high quality, with a value close to 1 for $H@3$ and even a $H@1$ between 0.61 and 0.91 - they clearly outperform a baseline random classifier. We also compute the overall performance over all primary emotion classes, as averages weighted by the number of instances corresponding to each class. The results are very good, with a value of 0.97 for $H@3$ and 0.93 for MRR. For primary emotions, correlation between class size and performance is medium: Pearson's r is 0.45 for $H@3$ and 0.63 for $H@1$, RP, and MRR. Thus, when misclassifying instances the classifier is biased to incorrectly assigning one of the two dominant classes "Fear" or "Sadness". The overall weighted results for the secondary human emotion label recommendations are almost identical with the case of primary emotions. Again, we find a classifier bias towards popular classes, correlation between *a priori* probability of a class and performance is smaller for secondary emotions (Pearson's r is between 0.32 and 0.37 for the different evaluation measures). Looking at the different mood classes individually, six out of 17 achieve 0.88 or higher for $H@1$ and 0.93 or more for $H@3$. For four moods, "Affection", "Zest", "Irritation" and "Lust", performance is considerably lower with $H@3$ ranging from 0.18 to 0.52. The main reason is the relatively small number of pictures contained in each of these groups, which made learning more difficult.

Compared to the music mood and theme recommendations, inference of moods/emotions for picture resources is of higher quality. This is probably due to the data which was used as ground truth: mood-related *Flickr* groups, manually created by users. The ground truth gathered from *AllMusic.com* had to be mapped to the hierarchy of human emotions to reduce the extremely high number of classes. This process potentially introduces some noise into the data. The results confirm once more the hypothesis on which we based our recommendation approach: existing tags can give good indications regarding the corresponding moods of the pictures/ moods and themes of music songs.

3. SEARCH AND RECOMMENDATIONS IN SOCIAL NETWORKS

From sociology as well as social network analysis as a discipline in its own right quite some descriptive statistics and generative models have become popular to characterize social networks. Well-known is that weak ties often act as bridges and thus hold potential for new information and the generation of creative ideas, job offers (e.g. [6, 10]). On the other hand, strong ties offer support and trust (e.g. see [15]) and show tendencies for homophily and transitivity [20, 9].

This means they likely share knowledge, preferences, values and friends. While many 'real-world' studies had to cope with design issues (amount of data, retrospective informant accuracy, etc.), online networks offer huge and interesting datasets to work with. With respect to exploiting online social connections for search and recommendation, two broad areas need to be studied:

- How can ties be modeled based on implicit and explicit indicators found in online social networks? How can these ties be characterized? Do we find homophily and heterophily as expected from offline social relations? How do these ties evolve over time?

- How useful are weak and strong ties in search and recommendation (information propagation)? How can they be incorporated to show their potential benefit in information systems (e.g. diversity/non-redundancy)? Are there restrictions by contexts or domains?

We address related work, open issues for both topics in the following sections to then conclude with planed future work.

3.1 Inferring and characterizing tie strength

In sociological theory, an impressive amount of work has been done regarding the measurement of tie strength. A lot of reliable indicators have been identified, e.g. interaction frequency, duration, intimacy (e.g. [18]), etc. ([8] for a quick overview). Kahanda and Neville [12] recently presented a machine learning approach to automatically identify strong friends. The authors formulated a link strength prediction task: For each friend pair (u,v) given their user profile attributes like age, gender, etc., their interactions (writing on the friend's wall, tagging a photo), and network information (e.g. number of mutual friends) a supervised learning method decides whether they are "top friends". Evaluation on data from the public Purdue Facebook network, where users can nominate best friends within the "Top friends" application, showed that with an AUC of 87% best friends can be successfully distinguised from weak ties. Those best classification results were achieved on network-transactional features (i.e., moderate transactional activity like wall posts by interactions with other users).Thus, user interactions are highly predictive, but it is also necessary to consider such transactional events in the context of user behavior within the larger social network. Surprisingly, attribute similarity features lead to classification results close to a random classifier, indicating that the homophily assumption does not hold for this Facebook network or the important attributes are not available in Facebook.

In a similar work, Gilbert and Karahalios [8] predict tie strength as a linear combination of 74 Facebook variables (e.g. last comment, num friends, wall words). Besides comparable results, a mapping of (sociological) dimensions to the different variables is provided enabling generalization of the approach. In [27], Yamamoto & Matsumura analyzed optimal heterophily between senders and receivers in terms of blogging influence (tracked via re-occurring terms and links) and domain knowledge. They found that the majority of pairs favor small heterophily, in particular people most often adopt topics or products when the sender is just slightly more influential.

3.2 Approaches exploiting ties

Approaches for efficiently searching and propagating information in online communities build strongly upon methods developed in social networks analysis. Real networks like the WWW, the Internet, spreading of diseases/epidemics, social, biological and linguistic networks have so far been studied mostly with respect to structure. Graph theoretic measures like density, indegrees, outdegrees, centrality, diameter, (structural) cohesion, etc. indicate the potentials and bottlenecks of a network. As an example, one experimental finding recently receiving a lot of attention is the 'small-world phenomenon' (also '6 degrees of separation')[14, 10]. While early works investigated patterns of communication in small, closed groups, recent work analyzed communication flows in huge social networks e.g. based on mobile phone calls [21], instant messaging [14], or the cascading propagation of information through the blogosphere [11]. In their studies, [21] found that weak ties are crucial for the structural integrity of the network. Strong ties, on the other hand, are important for sustaining local groups/cliques. Concerning information propagation both types of connections are not sufficient, the first due to infrequent, rare contact, the latter due to being bound to their local groups. Epidemic or gossip-based algorithms adapt such established patterns to enable efficient spread of information for distributed computing or in Peer-2-Peer systems [7].

Similarly, social search and recommendation algorithms try to exploit the communication and interaction patterns found in social networks as well as, for example, the trust and similarity typical for strong ties. Referral Web is a first approach to integrate social networks and Collaborative Filtering (CF) [13]. Here, a social network was constructed from cooccurring names in the WWW, for example, links on a home page or co-authorship. Queries that can be answered based on this network have the form "which connection do I have to XY" or "documents about databases by people close to XY". [26] models real-world information flows in order to give recommendations and rank users according to influence based on the usage of certain communications paths. For this, diffusion rate between users is computed based on access time/order to the same documents. The automatic evaluation shows that standard CF algorithms can be outperformed in accuracy by up to 80%. Moreover, the underlying social network can be used to overcome the sparsity/missing data problem, for example, by applying factor analysis on the user-item-matrix enriched with explicit user connections [16]. For personalized recommendations of new posts concerning some news item, [24] extend their CF recommender system in a way that strong social network ties (here: members of a thematic group) indicate a high value of a post with respect to completeness and simplicity. Weak ties, in contrast, imply diversity of opinions. From ratings given to posts the system learns a user's preference regarding completeness and diversity, to which recommendations are adapted. [23] presents a framework for social search and recommendation that integrates classical CF attributes for users and resources with an ontology and social connectivity (explicit friendship or 'spiritual', i.e. similar interests modeled via tags) within a scoring model. A small evaluation study shows that 'spiritual' connections in particular improve search results significantly, but not for all kinds of queries. Social query expansion by tags used by friends, however, did not lead to improved performance. In a related work, [1] demonstrated that social search, implemented as search among all friends having used a query term as tag before, possibly combined with an authority score for users can yield the best precision for search in *Flickr*. Also for efficiently searching inside collaborative tagging networks like *Delicious* incorporating social connections between users and between tags proved useful. A top-k algorithm combined with dynamic tag expansion and dynamically extending search over socially connected users can answer queries considerably faster than traditional approaches [22].

4. CONCLUSIONS & FUTURE WORK

Online social networks offer great data to analyze and experiment with for enhancing user profiling, search and recommendation. The concept of tie strength seems a promising framework for identifying the diverse potentials different online social relations can bring. First, collaborative tagging provides reliable, non-redundant and interpersonally valuable metadata, that can be used to enhance searchability of resources as well as estimate user-user or user-item similarities. For this, no explicit friendship relationships have to be given. The results of our comparative tagging analysis provide more insight into the use of different kinds of tags for improving search. With respect to weak and strong ties, more research is needed on how to model ties based on explicit or implicit (e.g. tags) indicators. We plan to conduct experiments on learning tie strength and exploiting it for search and recommendation within other kinds of social networks, e.g. the music platform *Last.fm*. For different domains, system designs and available transactional data, results may deviate from the previous studies. So far, there are still no unambiguous results regarding homophily in recent online networks. Characterizing the relationships people form online and studying how these relations (or their attributes like taste profile similarity) evolve over time are important to assess the value of ties for improving search and recommendations. Applying standard social network analysis procedures on the new large datasets will also shed additional light on larger community structures around strong and weak ties in general. More importantly, few work has been done so far on how to incorporate tie strength into information retrieval and recommendation. First experiments will use tie strength within the similarity computation of users (to users and items). Information diffusion within personal networks will be studied and a model derived on which social recommendations can be build.

5. ACKNOWLEDGMENTS

This work was partially supported by the GLOCAL project funded by the European Commission (Contract No. 248984).

6. REFERENCES

[1] M. Bender, T. Crecelius, M. Kacimi, S. Michel, T. Neumann, J. X. Parreira, R. Schenkel, and G. Weikum. Exploiting social relations for query expansion and result ranking. In *Proceedings ICDEW 2008*, 2008.

[2] K. Bischoff, C. S. Firan, W. Nejdl, and R. Paiu. Can all tags be used for search? In *Proceedings CIKM*, 2008.

[3] K. Bischoff, C. S. Firan, W. Nejdl, and R. Paiu. How do you feel about "dancing queen"?: deriving mood & theme annotations from user tags. In *Proceedings JCDL*, 2009.

[4] K. Bischoff, C. S. Firan, W. Nejdl, and R. Paiu. Bridging the gap between tagging and querying vocabularies: Analyses and applications for enhancing multimedia ir. *Journal of Web Semantics, Special Issue on "Bridging the Gap" - Data Mining and Social Network Analysis for Integrating Semantic Web and Web 2.0*, 2010.

[5] K. Bischoff, C. S. Firan, R. Paiu, W. Nejdl, C. Laurier, and M. Sordo. Music mood and theme classification - a hybrid approach. In *Proceedings ISMIR*, 2009.

[6] R. S. Burt. Structural holes and good ideas. *American Journal of Sociology*, 110:349–399, 2004.

[7] A. Demers, D. Greene, C. Hauser, W. Irish, J. Larson, S. Shenker, H. Sturgis, D. Swinehart, and D. Terry. Epidemic algorithms for replicated database maintenance. In *Proceedings PODC '87*, 1987.

[8] E. Gilbert and K. Karahalios. Predicting tie strength with social media. In *Proceedings CHI*, 2009.

[9] M. Granovetter. The strength of weak ties. *The American Journal of Sociology*, 78:1360–1380, 1973.

[10] M. Granovetter. The strength of weak ties: a network theory revisited. *Sociological Theory*, 1:201–233, 1983.

[11] D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins. Information diffusion through blogspace. In *Proceedings WWW*, 2004.

[12] I. Kahanda and J. Neville. Using transactional information to predict link strength in online social networks. In *Proceedings ICWSM*, 2009.

[13] H. Kautz, B. Selman, and M. Shah. Referral web: combining social networks and collaborative filtering. *Communications of the ACM*, 40(3):63–65, 1997.

[14] J. Leskovec and E. Horvitz. Planetary-scale views on a large instant-messaging network. In *Proceedings WWW*, 2008.

[15] D. Z. Levin and R. Cross. The strength of weak ties you can trust: The mediating role of trust in effective knowledge transfer. *Management Science*, 50(11):1477–1490, 2004.

[16] H. Ma, H. Yang, M. R. Lyu, and I. King. Sorec: social recommendation using probabilistic matrix factorization. In *Proceedings CIKM*, 2008.

[17] C. Marlow, M. Naaman, D. Boyd, and M. Davis. Ht06, tagging paper, taxonomy, flickr, academic article, to read. In *Proceedings HYPERTEXT*, 2006.

[18] P. V. Marsden and K. E. Campbell. Measuring tie strength. *Social Forces*, 63(2):482–501, 1984.

[19] A. McAfee. How to hit the enterprise 2.0 bullseye. online: http://andrewmcafee.org/2007/11/how_to_hit_the_enterprise_20_bullseye/, 2007.

[20] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27:415–444, 2001.

[21] J.-P. Onnela, J. Saramäki, J. Hyvönen, G. Szabó, D. Lazer, K. Kaski, J. Kertész, and A.-L. Barabási. Structure and tie strengths in mobile communication networks. *Proc. Natl. Acad. Sci.*, 104, 2008.

[22] R. Schenkel, T. Crecelius, M. Kacimi, S. Michel, T. Neumann, J. X. Parreira, and G. Weikum. Efficient top-k querying over social-tagging networks. In *Proceedings SIGIR*, 2008.

[23] R. Schenkel, T. Crecelius, M. Kacimi, T. Neumann, J. X. Parreira, M. Spaniol, and G. Weikum. Social wisdom for search and recommendation. *IEEE Data Eng. Bull.*, 31(2):40–49, 2008.

[24] A. Seth and J. Zhang. A social network based approach to personalized recommendation of participatory media content. In *Proceedings ICWSM*, 2008.

[25] P. Shaver, J. Schwartz, D. Kirson, and C. O'Connor. Emotion knowledge: Further exploration of a prototype approach. *Journal of Personality and Social Psychology*, 52(6):1061–1086, 1987.

[26] X. Song, Y. Chi, K. Hino, and B. L. Tseng. Information flow modeling based on diffusion rate for prediction and ranking. In *Proceedings WWW*, 2007.

[27] H. Yamamoto and N. Matsumura. Optimal heterophily for word-of-mouth diffusion. In *Proceedings ICWSM*, 2009.