INFERRING EMOTIONS FROM IMAGE SOCIAL NETWORKS USING GROUP-BASED FACTOR GRAPH MODEL

Wenjing Cai[1,2,3], *Jia Jia*[1,2,3], *Wentao Han*[1*]

[1]Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China
[2]Key Laboratory of Pervasive Computing, Ministry of Education
[3]Beijing National Research Center for Information Science and Technology
caiwj425@163.com,jjia@mail.tsinghua.edu.cn,hanwentao@tsinghua.edu.cn

ABSTRACT

Inferring emotions from image social networks is a hot research topic nowadays. For image social networks (Flickr, Instagram), there is an interesting phenomenon that people would like to establish or attend virtual groups and share images with different topics and emotions in different groups. Previous researches on inferring emotions usually focus on image content and user personalization, thus leading an interesting but challenging problem: whether virtual groups can influence members(users)' emotions. In this paper, we systematically study this problem from two aspects: 1) whether *group homophily in users' emotions* exists in image social networks; 2) how to model this subtle and complex group homophily in image social networks. Inspired by the study results of two aspects, we introduce group information to infer emotions in image social networks, and propose a novel *Group-Based Factor Graph Model (G-FGM)*, incorporating image content, user personalization and group information to understand the emotions behind social images better. The experimental results on a dataset containing 218,816 emotion-labeled images from Flickr show that our model outperforms (8.6-19.4% improvement in terms of F1-Measure) several baseline methods.

***Index Terms*—** social networks, image, groups, emotion inference

1. INTRODUCTION

With the development of social networks, more and more people are willing to share images to express their feelings via social network platforms such as Flickr[1] and Instagram[2]. For the world's largest image-sharing website Flickr, 38% of images are explicitly annotated as positive or negative emotion tags by their publishers [1]. Understanding the emotions behind images in social networks is an important research topic with big challenges. It can also benefit many applications such as image retrieval, market advertising and personalized recommendation.

At present, the studies on inferring emotions from social images mainly focus on the image content and user personalization. Wang [2] proposed the interpretable aesthetic features inspired by art theories. [3] developed a SentiBank to detect the sentiments reflected in image visual content using the Adjective Noun Pair (ANP) concept. And [4] extracted the principles-of-art-based emotion features (PAEF) for affective image classification and regression. Other types of mid-level attributes of image content can also be found like Sentribute [5]. Additionally, inferring emotions is also influenced by user personalization as different users express emotions differently because of different social and cultural backgrounds [6]. Specifically, users' demographics have an effect on the emotion expressions of users [1]. And, the users' emotions at present are also influenced by their past emotions [7] and others' emotions by image sharing [7],[8].

Besides image content and user personalization, there is an interesting phenomenon that people would like to establish or attend virtual groups and share images with different topics and emotions in different groups. Groups provide a platform where users can communicate and work for the same topics together. Our statistics show that in Flickr there are a total of 466,099 groups including 1,255,478 users joined before December 2013. [9] summed up four factors that affect users when they join a group, indicating similar users tend to join similar groups. What's more, it's found that two users who communicate frequently are more likely to join some similar groups [10]. From the motivation of users joining a group, most users within a group have similar parts, such as interests, viewpoints, etc, which can be called *group homophily*. Inspired by this, we propose two basic questions: 1)whether *group homophily in users' emotions* exists in image social networks; 2)how to model this subtle and complex group homophily in image social networks.

In this paper, we first explore the above two basic questions through data observations which demonstrate the existence of *group homophily in users' emotions* and inspire us to introduce group information to model this group homophily.

[*]Corresponding author: Wentao Han(hanwentao@tsinghua.edu.cn)
[1]http://www.flickr.com
[2]https://www.instagram.com

Based on the observations, we propose a novel model named group-based factor graph model (G-FGM) to infer emotions in image social networks, leveraging the features from three aspects: image content, user personalization and group information. The experimental results on a dataset containing 218,816 emotion-labeled images from Flickr show that our model outperform (8.6-19.4% improvement in terms of F1-Measure) several baseline methods and also prove that groups actually influence members(users)' emotions. We also conduct some case studies to further demonstrate that the group influence is weakening with a longer time interval and this influence on negative emotions is weakening faster than positive emotions.

2. PROBLEM FORMULATION

Image social network: Given a time-varying image social network G, we have $G = \{V, E^t, X, G_I\}$. Specifically, V is the set of users who publish the images, $V = \{v_i\}$; each user v_i is instantiated as his(her) demographic $\mathbf{d}_i$ which includes the user's information about gender, occupation, marital status [1] and location [6]. E^t presents the friendships among users at time t. $X = \{x_i^{j,t}\}$, is the images set. $x_i^{j,t}$, the jth image published by user v_i at time t, is instantiated as the visual features $\mathbf{u}_i^j$ adopt from [2]. $G_I = \{\mathbf{g}_i^j\}$ is the set of group information. $\mathbf{g}_i^j$ is the information of the group where image $x_i^{j,t}$ is published.

Group information $\mathbf{g}_i^j$: $\mathbf{g}_i^j = \langle$*group emotion*, *group size*, *group ratio of social role*, *group connectivity*$\rangle$. *Group emotion* is the main emotion of users who publish images in this group. *Group size* refers to the number of all users in a group. *Group ratio of social role* consists of *opinion leader ratio* and *structural hole spanner ratio*, where *opinion leader(structural hole spanner) ratio* refers to the ratio of users who are opinion leaders(structural hole spanners) [11] in a group. *Group connectivity* presents the connectivity of a group where users are vertexes and friendships are edges. Given a group g, which has n_g users and m_g friendships, formal definition is followed as:

$$group\_connectivity = \frac{\sum_{l=1}^{l=m_g} e_l * w_l}{\sum_{l=1}^{l=2\binom{n}{2}} e_l * w_l} \tag{1}$$

where e_l is the lth edge which has a weight w_l obtained from the number of comments among vertexes (users). Specifically, the edge e_l is 1, if this edge exists, otherwise, it is 0.

Emotion: y_i^t indicates the emotion of user v_i at time t. we have an key intuition that the users' emotions are expressed by their published images. And, based on the Ekman's emotion theory [12], we adopt the emotion space $S = \{happiness, surprise, anger, disgust, fear, sadness\}$.

Problem: Given a time-varying image social network G, $G = \{V, E^t, X, G_I\}$, our goal is to learn a function f using labeled data to predict the users' emotions expressed by emotion-unlabeled images.

$$f : G = \{V, E^t, X, G_I\} \Rightarrow Y \tag{2}$$

where $Y = \{y_i^t\}$, and $y_i^t \in S$.

Table 1. The Statistical Results of PGE

	Average	Happy	Surprise	Anger	Disgust	Fear	Sad
Mean	0.611	0.624	0.502	0.623	0.649	0.553	0.629
Var	0.031	0.031	0.023	0.041	0.042	0.019	0.036

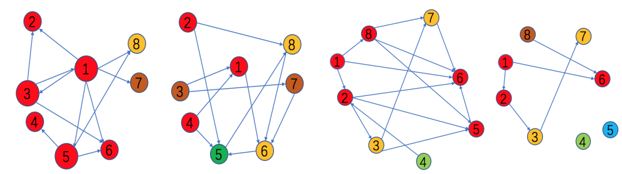

(a) groups with high and low ratio of opinion leaders (b) groups with high and low connectivity

Fig. 1. Graphical representation of different groups. The node is a user; the line with an arrow is the friendships among users; different colors on the nodes represent different emotions of users. And the size of a node present the social role of a user.

3. DATA OBSERVATION

In this section, we conduct several data observations to explore the above problem from two aspects: 1) whether *group homophily in users' emotions*(GHE for short) exists in image social networks; 2)how to model this subtle and complex group homophily in image social networks.

3.1. Data Set

To prepare our observations, we randomly download 2,060,353 images which are published by 1,255,478 users in 466,099 groups from Flickr. Before observations, we need to label these images with emotion categories. Facing the plenty of images, it is impossible to label these images manually. We adopt an automatic labeling method, which has been used in previous works [1],[13],[7]. Herein we first construct word lists for each emotion category using WordNet[3] and HowNet[4]. Then, we compare the images' tags provided by publishers and comments with these word lists. Finally, we label these images with one emotion category whose word list matches tags and comments most. After automatic labeling, we get 354,192 emotion-labeled images from 69,430 groups.

3.2. Observation on *group homophily in users' emotions*

We choose two weeks as a time slice, and calculate the *proportion of groups' main emotion*(PGE) in each time slice. In order to distinguish six emotion categories, we also consider the PGE in particular groups whose main emotions are the six categories respectively. For instance, 'Happy' presents the groups whose main emotions are *happiness* category, similarly for the other five emotion categories. Table 1 lists the statistical results of PGE. It can be seen that the mean of PGE is between 0.5 and 0.7, and the variance(Var)[5] is about 0.03, indicating the users within a group are more likely to

[3]http://wordnet.princeton.edu/
[4]http://www.keenage.com/
[5]https://en.wikipedia.org/wiki/Variance

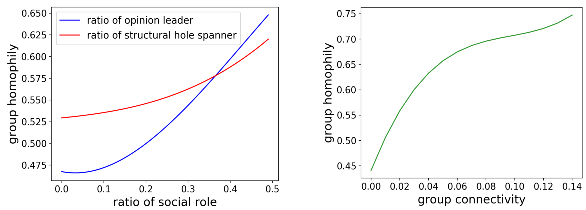

(a) GHE and ratio of social role. (b) GHE and connectivity

Fig. 2. Relationship between GHE and group information.

express the same emotion, which demonstrates the existence of GHE in image social networks.

3.3. Observation on the influenced features on GHE

After proving the existence of GHE, we are considering how to model GHE in image social networks. We explore the relationships between GHE and influenced features extracted from the attributes of users and their friendships (defined as group information in Section 2). The graphical representation of different groups lists in Figure 1.

Relationship between group homophily and group ratio of social role. In our work, we use *PageRank* [14] to determine whether a user is an opinion leader, and we use *Network Constraint Score* [15] to determine whether a user is a structural hole spanner. According to the result in Figure 2(a), we can find that with the increase in group ratio of social role (opinion leader or structural hole spanner), the group homophily is also increasing. That is, there is a positive effect of group ratio of social role on GHE.

Relationship between group homophily and group connectivity. We calculate the connectivity of each group and get the relationship between GHE and connectivity. The result in Figure 2(b) shows that there is a positive effect of group connectivity on GHE.

The summary of observations is as follows:

- The *group homophily in users' emotions* exists in image social networks.
- For different groups, the *group homophily in users' emotions* differs. We introduce *group ratio of social role* which describes the features of group members and *group connectivity* which describes the group's holistic features to model this subtle and complex group homophily in image social networks.

4. MODEL

Inspired by the graph structure groups, we propose a novel unified model, G-FGM, to infer emotions leveraging several factors in three aspects. We define the factors which affect emotion inference as factor functions in the factor graph model. The objective function is defined based on the joint probability of factor functions [1],[16]. The learning task of our model is to maximize the joint probability by parameters adjustment.

4.1. Factors definition.

The definitions and instantiations of factors in our model are as follows:

- **Content factor** $f_1(\mathbf{u}_i^j, y_i^t)$: It presents how the user's emotion y_i^t is induced by visual features of published image $x_i^{j,t}$.
$$f_1(\mathbf{u}_i^j, y_i^t) = \frac{1}{Z_\alpha}\exp\{\alpha^T \cdot \mathbf{u}_i^j\} \tag{3}$$
- **User personalization correlation**: This correlation presents the effect of user personalization on emotion reference. It has three factors: temporal correlation factor$f_2(y_i^{t'}, y_i^t)$ presents the correlation between a user' emotions at times t' and t, where $t' < t$;
$$f_2(y_i^{t'}, y_i^t) = \frac{1}{Z_\varepsilon}\exp\{\varepsilon_i \cdot H(y_i^{t'}, y_i^t)\} \tag{4}$$
user's demographic factor$f_3(\mathbf{d}_i, y_i^t)$ presents the effect of user's demographic on user's emotion y_i^t;
$$f_3(\mathbf{d}_i, y_i^t) = \frac{1}{Z_\delta}\exp\{\delta^T \cdot \mathbf{d}_i\} \tag{5}$$
friendship correlation factor$f_4(y_{i_1}^t, y_i^t)$ depicts the emotional influence among users.
$$f_4(y_{i_1}^t, y_i^t) = \frac{1}{Z_\eta}\exp\{\eta_{i_1,i} \cdot H(y_{i_1}^t, y_i^t\} \tag{6}$$
where $\eta_{i_1,i}$ presents the influence weight between user v_{i_1} and v_i.
- **Group information factor**$f_5(\mathbf{g}_i^j, y_i^t)$: This factor describes the correlation between the group information and it's member's emotions.
$$f_5(\mathbf{g}_i^j, y_i^t) = \frac{1}{Z_\lambda}\exp\{\lambda^T \cdot \mathbf{g}_i^j\} \tag{7}$$
As for the above functions, α, ε, η, δ and λ are the training parameters of the factor graph model. Z_α, Z_ε, Z_η, Z_δ, Z_λ are normalization terms, and function H is defined as a vector of indicator functions.

4.2. Model learning

Given the above definitions, the joint distribution of our model is:

$$P(Y|G) = \prod_{x_i^{j,t}} f_1(\mathbf{u}_i^j, y_i^t) \prod_{x_i^{j,t}} \prod_{y_i^{t'}} f_2(y_i^{t'}, y_i^t)$$
$$\prod_{x_i^{j,t}} f_3(\mathbf{d}_i, y_i^t) \prod_{x_i^{j,t}} \prod_{y_{i_1}^t} f_4(y_{i_1}^t, y_i^t) \prod_{x_i^{j,t}} f_5(\mathbf{g}_i^j, y_i^t) \tag{8}$$

$$P(Y|G) = \frac{1}{Z}\exp\{\sum_{x_i^{j,t}} \alpha^T \mathbf{u}_i^j\}$$
$$\times \exp\{\sum_{x_i^{j,t}} \sum_{y_i^{t'}} \varepsilon_i H(y_i^{t'}, y_i^t)\} \times \exp\{\sum_{x_i^{j,t}} \delta^T \mathbf{d}_i\}$$
$$\times \exp\{\sum_{x_i^{j,t}} \sum_{y_{i_1}^t} \eta_{i_1,i} H(y_{i_1}^t, y_i^t)\} \times \exp\{\sum_{x_i^{j,t}} \lambda^T \mathbf{g}_i^j\} \tag{9}$$

And the log-likehood objective function is
$$\Psi = \log P(Y|G) = \log \sum_{Y|Y^U} \exp\{\theta^T K\} - \log Z \tag{10}$$
where Z is a normalization term; and $\theta = \{\alpha, \varepsilon_i, \eta_{i_1,i}, \delta, \lambda\}$.

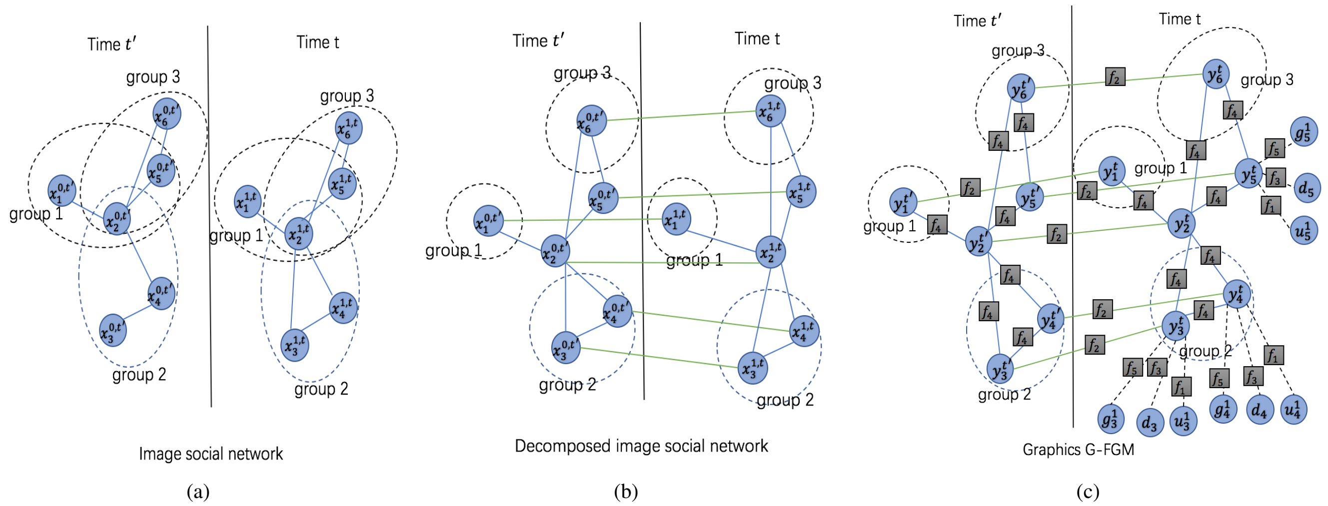

Fig. 3. Graphical representation of G-FGM. Blue edge presents the friendships among users who publish the images, green edge presents the temporal correlation. **Note that $x_5^{1,t}$ is published both in group 1 and group 3, and it's group information $\mathbf{g}_5^1$ is related to the information of group 1 and group 3**. In image social network, there are some images published in several different groups simultaneously, such as $x_2^{0,t'}$ and $x_5^{0,t'}$. In this case, we separate these images from others, such as in (b). The group information of these images is obtained from all published groups, like $\mathbf{g}_5^1$ in (c). And for space constraints, in (c), we only display part of these images with all factors.

K aggregates all the factor functions over all nodes, namely, $K = \sum_{x_i^{j,t}} k(y_i)$, where $k(y_i) = \{\mathbf{u}_i^j, H(y_i^{t'}, y_i^t), \mathbf{d}_i, H(y_{i_1}^t, y_i^t), \mathbf{g}_i^j\}$

Training the model is to maximize the function Ψ, and the gradient of θ can be calculated as:

$$\begin{aligned}\frac{\partial \Psi}{\partial \theta} &= \frac{\partial \log P(Y|G)}{\partial \theta} \\ &= E_{p_\theta(Y|Y^U,G)}K - E_{p_\theta(Y|G)}K\end{aligned} \tag{11}$$

To get the gradient of θ, we calculate the values of $E_{p_\theta(Y|Y^U,G)}K$ and $E_{p_\theta(Y|G)}K$ using LBP(Loopy Belief Propagation). Then, we update θ by $\theta = \theta_0 + \kappa \cdot \frac{\partial \Psi}{\partial \theta}$ where θ_0 is the initial value of θ, and κ is the learning rate.

5. EXPERIMENTS

5.1. Experimental set up

Data set We use data from the dataset in Section 3. As our model is a time-varying model, we have to use images with published time to conduct experiments. After removing images without time, we get 218,816 images which contain 46.2% *happiness-labeled* images, 9.7% *surprise-labeled* images, 8.0% *anger-labeled* images, 5.3% *disgust-labeled* images, 17.3% *fear-labeled* images, 13.5% *sadness-labeled* images. As we can see, the numbers of six emotion categories are imbalanced, but the numbers of positive(101,189) and negative(117,627) emotions are almost balanced, which is quite practical and reasonable. Facing the fact that imbalanced data badly hurts the performance, we use SMOTE [17] to do oversampling for the training data before classification.

Comparison methods. In order to accurately and clearly see the effectiveness of our model, we carry out some comparative experiments. We use Support Vector Machine (SVM) [18], Deep Neural Network(DNN) [19] and D-FGM [1] for comparison.

Metrics. In our experiments, we calculate precision, recall and F1-Measure[6] to evaluate the performance of methods. And the 5-fold cross validation is used in our experiments.

5.2. Results and analyses

Prediction performance. Table 2 lists the results of our model and other methods. We can see that G-FGM achieves a best performance than other methods. In terms of F1-Measure, our model reaches 50.4% in average, improved by 8.6% compared with D-FGM, 19.0% compared with DNN, 19.4% compared with SVM, confirming the effectiveness of our model.

DNN and SVM have no ability to handle the correlation features(such as temporal correlation and friendship correlation), which hurts the performance. D-FGM ignores the group information that is important in image social networks. Our model combines not only image content and user personalization but also the group information, which makes the emotion prediction more accurate and effective. Importantly, the result shows that both D-FGM and G-FGM have a better performance than DNN, indicating the factor graph model has more advantages in modeling social networks. Besides, although SMOTE has improved the bad effect of imbalanced data, the performance for minority classes (”surprise” and ”anger”) is still poor.

[6] https://en.wikipedia.org/wiki/F1_score

Table 2. The result of comparative experiments

Method	Happiness	Surprise	Anger	Disgust	Fear	Sadness	Average
SVM	0.640	0.133	0.134	0.386	0.222	0.343	0.310
DNN	0.400	0.207	0.234	0.348	0.350	0.343	0.314
D-FGM	0.739	0.223	0.170	0.270	0.623	0.482	0.418
G-FGM	**0.760**	**0.256**	**0.297**	**0.474**	**0.677**	**0.559**	**0.504**

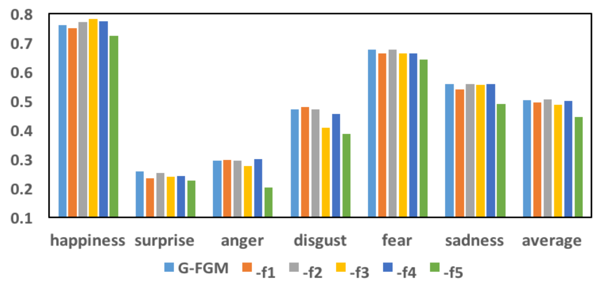

Fig. 4. F1-Measure of different factor contribution.

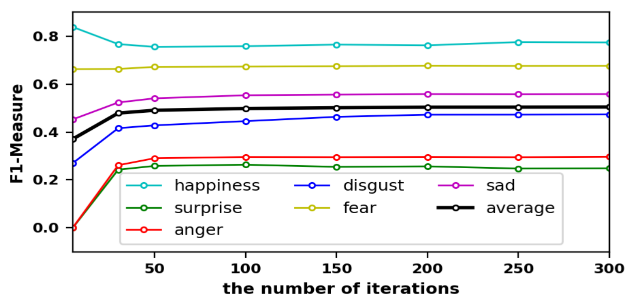

Fig. 5. F1-Measure of different number of iterations.

Factor contribution analysis. Here, we show the contribution of each factor to the proposed model. Specifically, we eliminate each factor from the model in turn, and compare the performance of emotion inference. The results are shown in Figure 4. We can find that for most emotion categories, each factor contributes to the model, which demonstrates the effectiveness of each factor especially the group information factor. And we also have some other interesting findings, which are summarized as followed:

- As for *anger* and *disgust*, group information factor(f_5) contributes most to the performance of our model (+9.2% for *anger* and +8.5% for *disgust*).
- User's demographic factor f_3 has contribution to most emotion categories, while it has little contribution to *happiness*, indicating *happiness* is much more common for people, while *disgust* is much more affected by personal information.
- Temporal correlation factor f_2 and friendship correlation factor f_4 have a small contribution for most emotions, as the oversampled nodes have no explicit edge relationships (f_2, f_4).

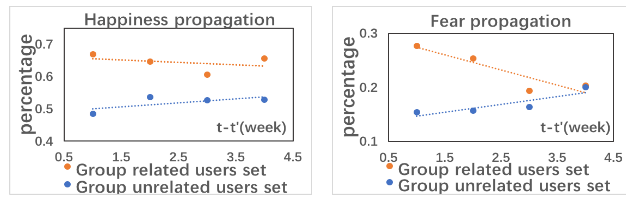

Fig. 6. Case Study

Parameter sensitivity analysis. Here we show parameters sensitivity analysis about *the number of iterations*. We evaluate the performance of the proposed model as the number of iterations increases. Figure 5 lists the result of each emotion category in terms of F1-measure. When the number of iterations is too small to train the model, the prediction result tend to be *happiness* category, because the number of *happiness-labeled* images is the largest. As the number of iterations increases, the average F1-measure increases accordingly and finally converges.

5.3. Case study.

In the end, we would like to show some interesting cases to study how groups affect users' emotions.

Happiness propagation from groups to users. In order to verify the happiness propagation from groups to users, we conduct a sampling test [7]. Specifically, as for *happiness*, we define two sets of users: $S_{groupRelated}$ who publish images at time t and have groups whose main emotion is *happiness* at time t'; $S_{groupUnrelated}$ who publish images at time t and have no groups whose main emotion is *happiness* at time t' ($t' < t$). For each set of users, we calculate the percentage of *happiness* people varying time interval $t - t'$. Figure 6(a) lists the result, where we can see that the overall percentage of $S_{groupRelated}$ is higher than that of $S_{groupUnrelated}$, confirming users in $S_{groupRelated}$ are more likely to be happy with the influence of the group 'happiness'. And the downward trend line shows the influence is weakening with a longer time interval.

Fear propagation from groups to users. We further use the same method in the first case study to see the fear propagation from groups to users. The result in Figure 6(b) shows that users in $S_{groupRelated}$ are more likely to fear with the influence of the group 'fear'. And the downward trend line shows the influence is weakening with a longer time interval.

Comparing the results of two case studies above, we also find that the trend line of negative emotion (*fear*) is steeper than that of positive emotion (*happiness*), indicating the group influence on negative emotion is weakening faster over time.

6. CONCLUSION

In this paper, by introducing group information, we implement a joint emotion inference model combining image content, user personalization and group information. Our experiments prove that groups influence members' emotions and help to improve inferring emotions from image social networks. We can further study the evolution of groups and the emotion inference using dynamic groups in image social networks.

7. ACKNOWLEDGEMENTS

This work is supported by National Key Research and Development Plan (2016YFB1001200),the Innovation Method Fund of China (2016IM010200),the National Natural, and Science Foundation of China (61521002).

8. REFERENCES

[1] Boya Wu, Jia Jia, and Yang Yang, "Understanding the emotions behind social images: Inferring with user demographics," in *IEEE International Conference on Multimedia and Expo*, 2015, pp. 1–6.

[2] Xiaohui Wang, Jia Jia, Jiaming Yin, and Lianhong Cai, "Interpretable aesthetic features for affective image classification," in *IEEE International Conference on Image Processing*, 2014, pp. 3230–3234.

[3] Damian Borth, Rongrong Ji, Tao Chen, Thomas Breuel, and Shih Fu Chang, "Large-scale visual sentiment ontology and detectors using adjective noun pairs," in *ACM International Conference on Multimedia*, 2013, pp. 223–232.

[4] Sicheng Zhao, Yue Gao, Xiaolei Jiang, Hongxun Yao, Tat Seng Chua, and Xiaoshuai Sun, "Exploring principles-of-art features for image emotion recognition," pp. 47–56, 2014.

[5] Jianbo Yuan, Sean Mcdonough, Quanzeng You, and Jiebo Luo, "Sentribute:image sentiment analysis from a mid-level perspective," in *International Workshop on Issues of Sentiment Discovery and Opinion Mining*, 2013, pp. 1–8.

[6] Sicheng Zhao, Hongxun Yao, Yue Gao, Rongrong Ji, Wenlong Xie, Xiaolei Jiang, and Tat Seng Chua, "Predicting personalized emotion perceptions of social images," in *ACM on Multimedia Conference*, 2016, pp. 1385–1394.

[7] Xiaohui Wang, Jia Jia, Jie Tang, Boya Wu, and Lianhong Cai, "Modeling emotion influence in image social networks," *IEEE Transactions on Affective Computing*, pp. 286–297, 2015.

[8] Yang Yang, Jia Jia, Shumei Zhang, Boya Wu, and Qicong Chen, "How do your friends on social media disclose your emotions?," in *Twenty-Eighth AAAI Conference on Artificial Intelligence*, 2014, pp. 306–312.

[9] Xiaolin Shi, Jun Zhu, Rui Cai, and Lei Zhang, "User grouping behavior in online forums," in *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Paris, France, June 28 - July*, 2009, pp. 777–786.

[10] Jimeng Sun and Jie Tang, *A Survey of Models and Algorithms for Social Influence Analysis*, Springer US, 2011.

[11] Yang Yang, Jie Tang, Cane Wing-ki Leung, and Yizhou Sun, "Rain: Social role-aware information diffusion.," in *AAAI*, 2015, pp. 367–373.

[12] Paul Ekman, "An argument for basic emotions.," *Cognition & Emotion*, pp. 169–200, 1992.

[13] Yang Yang, Jia Jia, Boya Wu, and Jie Tang, "Social role-aware emotion contagion in image social networks," in *AAAI Conference on Artificial Intelligence*, 2016.

[14] L Page, "The pagerank citation ranking : Bringing order to the web," *Stanford Digital Libraries Working Paper*, pp. 1–14, 1999.

[15] Burt and Ronald, "Structural holes versus network closure as social capital," *Nan*, 2001.

[16] Jie Tang, Yuan Zhang, Jimeng Sun, Jinhai Rao, Wenjing Yu, Yiran Chen, and A. C. M Fong, "Quantitative study of individual emotional states in social networks," *IEEE Transactions on Affective Computing*, pp. 132–144, 2012.

[17] Guillaume Lemaître, Fernando Nogueira, and Christos K. Aridas, "Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning," *Journal of Machine Learning Research*, 2017.

[18] Chih Chung Chang and Chih Jen Lin, "Libsvm: A library for support vector machines," *Acm Transactions on Intelligent Systems & Technology*, vol. 2, no. 3, pp. 1–27, 2011.

[19] Mariette Awad and Rahul Khanna, *Deep Neural Networks*, Apress, 2015.