

Learning to Predict Image Affects in Social Networks

Sen Wu
Department of Computer Science
Stanford University
senwu@stanford.edu

Zhihao Jia
Department of Computer Science
Stanford University
zhihao@stanford.edu

1. INTRODUCTION

Popular social networks such as Flickr, Facebook, Google+, and Twitter provide users ability to upload images and label them with certain emotions. For example, if you feel happy, you can upload a California sunshine picture on your Facebook and marked it with a smile face. Or you could feel grieved and update a heavy downpour picture on your Flickr with a sad emotion marker. Since images provide an efficient way to express emotions and affect others, its applications in social networks is becoming more and more popular. However, predicting image affects in social networks is a difficult problem and poses many challenges. The main **challenges** are:

- How to address the issue caused by the fact that similar images may cause totally different affects?
- How to model the correlation between the image visual features and affects?
- How to come up with a practical learning model and deploy it in a efficient way?

A naive approach to predict the user emotions by images is scanning basic elements in the images and exploit the correlations between certain image components and certain user emotions. For example, Sunshine usually correlates with happy emotion and downpour correlates with sad emotion. Another example is that images with warm colors correlates with positive affects, while images with cold colors correlates with negative affects. Therefore, by exploiting particular components (e.g., color features, brightness, and saturation) in images, we can predict the affects for some images.

One downside of the above approach is that similar images may cause totally different affects. Images with a need fire may refer happy to some people, but could mean angry to others. Therefore, how to predict the affects for these seemingly ambiguous images is a main challenge for our project. To address the issues, we introduce homophily phenomenon into our predictor, and refine our prediction by considering the affects homophily among users in the same community.

The “homophily phenomenon”, the observed tendency of “individuals tends to choose friends with similar tastes and preferences”[9], is one of the most striking and interesting regularities of social life. Homophily was first studied by psychologists with small groups of participants. With the rapid proliferation of online social networks such as Twitter, Facebook and Flickr, it becomes feasible to conduct investigations on real large-scale social networks [12]. In practice, it can help us achieve an in-depth understanding on the complex dynamics in social networks, and benefit various aspects such as advertising and economics.

In this paper, we employ an image-based social network Flickr as the source for our experimental data. Homophily can be influenced by various factors such as individual status and community pressure. In individual level, different tastes and preferences may

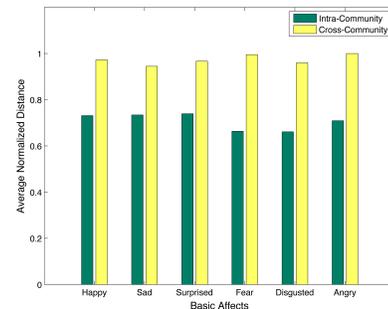


Figure 1: Community homophily analysis. Taking “Night Images” community (one of the most popular communities in Flickr) as the example, the figure shows the color features’ consistency within community and between community. The within-community average distance stands for the average euclidean distance between two images’ color features over all possible pairs coming from the same community. The between-community average distance stands for the average euclidean distance over all possible pairs coming from different communities.

result in different user emotions under a particular image [7, 15]; but in community level, connected individuals can influence each other and tend to behave similarly [1]. For example, as the common sense, colors are the most relevant attributes with image affects. Figure 1 shows the consistency of color features among images from the same or different communities in Flickr. It can be clearly seen that for all the six basic affects [5], the colors of their corresponding images in the same community have much higher consistency than those in different ones, which indicates the importance of homophily in predicting image affects in community-structured social networks.

How to design a principled model to automatically predict emotions from social images by considering both the image attributes (e.g., color features) and the homophily phenomenon (e.g., community information)? This problem is a non-trivial and pose a challenge. To address this challenge, we first formally define two patterns of homophily phenomenon in different levels: individual level and community level. A novel community-aware factor graph model (ComFG) is further proposed to incorporate both the image attributes and homophily patterns to predict affects of social images. To validate the effectiveness of our method, we test the proposed method on 150,000 random-download Flickr images, and the experimental results show that the proposed method significantly outperforms the alternative methods.

The rest of paper is organized as follows. In Section 2, we give the definition of homophily patterns and formulate the problem. Section 3 introduces the community-aware semi-supervised factor graph model and the learning algorithms. Then in Section 4, we report the experiment setup and performance using proposed method,

and further provides discussions. In Section 5, we review the related work. The paper concludes in Section 6.

2. PROBLEM DEFINITION

Let $G = (V, E)$ represent a social image network, where $V = \{p_1, \dots, p_n\}$ is the set of $|V| = n$ images with user u_i published p_i at time t_i and $E \subset V \times V$ is the set of $|E| = m$ edges. Each edge e_{ij} represents image p_i having a correlation with image p_j (e.g. p_i and p_j uploaded by the same user in a short time). In this network, we further assume all images belong to c communities and use an binary-valued $n \times c$ matrix \mathbf{C} to represent users' community memberships where each binary-valued element $c_{ik} \in C$ represents whether image p_i belongs to the k^{th} community. Here, we use C_k to represent k^{th} community. For example, in Flickr, users can construct communities and upload images to communities. In addition, each image v_i is associated with a d dimension attributes \mathbf{x}_i which describe images' attributes (e.g. color features). Here, \mathbf{X} is defined a $|V|^2 \times d$ attribute matrix associated with all images. Furthermore, each image uploaded by user can present his/her affects. A (a, p, u_i, t) denotes user u_i uploaded image $p \in \mathcal{P}$ at time t , where \mathcal{P} is the set of images, presents the affect $a \in \mathcal{A}$, where \mathcal{A} is the set of possible affects.

Our goal is to study how homophily behavior effect users' affects. Given this, we define several definitions and our problem.

Definition 1. Affects: The affective categories of an image p_i is denoted as $a_i \subset \mathcal{A}$, where \mathcal{A} is the affective space defined by Ekman[5].

$$\mathcal{A} = \{happy, sad, anger, disgust, surprise, fear\} \quad (1)$$

In order to better understanding the problem, we define two different levels of homophily from the level of individual and community. Please note that a image may involve in many different communities, here we only consider its main group.

Definition 2. Individual Homophily: The individual homophily is defined as the color attributes absolute error between image p_i and the mean color attributes from user u_j .

$$ih(p) = \frac{|a_p - a_{u_j}|}{a_{u_j}} \quad (2)$$

where a_p denotes the color attributes of image p , and a_{u_j} represents the mean color attributes of images uploaded by user u_j .

We also define community homophily to represent how communities effect the image affect.

Definition 3. Community Homophily: The community homophily is defined as the color attributes absolute error between image p_i and the community C_k mean color attributes image p_i belongs to.

$$gh(p, C_{pk}) = \frac{|a_{p_i} - a_{C_k}|}{a_{C_k}} \quad (3)$$

where a_{p_i} denotes the color attributes of image p_i , and a_{C_k} represents the community C_k 's mean color attributes.

Please note that individual and community homophily is defined on each specific affect. More precisely, the problem can be defined as:

Problem 1. Image Affect Prediction. Given a network $G = (V, E, \mathbf{C}, \mathbf{X})$ and an pictures uploading history $H = \{(a, p, v_i, t)\}_{p,a,i,t}$, our goal is to incorporate all kinds of information into a unified model to learn a predictive function:

$$f : \{G, H\} \rightarrow \mathcal{A} \quad (4)$$

where $a_i \in \mathcal{A}$ indicates what is the affect of the picture p_i .

3. MODEL FRAMEWORK

In this section, we propose a novel Community-aware Factor Graph (ComFG) model to incorporate both color features and homophily behaviour for better modeling and inferring the images' affects updated by users in large networks.

3.1 Community-aware Factor Graph Model

A social image network can be represented by a graph $G = (V, E)$ with a node set $|V| = n$ and an edge set $|E| = m$ representing all images and edges in network. Besides these, we know that there are many communities(C) in social networks and there are some special affects representations within communities. Figure 2 demonstrates the graphical representation of the ComFG model. Figure 2(a) shows our input is a social network with image attributes, users' relationship and community informations. Then, we split the network based on community information as Figure 2(b). Next, we define factor graph model on each community.

In ComFG model, we attempt to maximize conditional probability of image affects given their the corresponding attributes, communities and network, i.e., $P(Y|\mathbf{X}, G)$ and according to the Bayes' theorem, we can define a factor graph as,

$$P(Y|\mathbf{X}, G) = \frac{P(\mathbf{X}, G|Y)P(Y)}{P(\mathbf{X}, G)} \propto P(\mathbf{X}|Y) \cdot P(Y|G) \quad (5)$$

where $P(Y|G)$ defines the probability of affects given the structure fo the network and $P(\mathbf{X}|Y)$ defines the probability of generating the observation variables \mathbf{X} given the latent variables Y . Further assuming, the generating probability of the observation variables is conditionally independent, so

$$P(Y|\mathbf{X}, G) \propto P(Y|G) \prod_i P(\mathbf{x}_i|y_i) \quad (6)$$

where $P(\mathbf{x}_i|y_i)$ is the probability of generating the observation variables \mathbf{x}_i given the latent variable y_i .

In the proposed ComFG, we are capturing two kinds of information, i.e, the attributes associated with each image and two levels of homophily we defined in §2. We have two factor functions to represent the individual homophily and community homophily.

- **Individual homophily factor:** $f(y_i, ih(p_i))$ represents the correlation between image p_i 's affect and individual homophily.
- **Community homophily factor:** $f(y_i, gh(p_i, C_k))$ represents the correlation between image p_i 's affect and homophily to its community C_k .

More precisely, we decompose the original factor graph model into several communities based on the community information. As we know, there are c communities in the network, so the prior probability $P(\mathbf{x}_i|y_i)$ can be split into c communities as $P_k(\mathbf{x}_i|y_i)$. Then we can get,

$$P(Y|\mathbf{X}, G) \propto \prod_{k=1}^c P_k(Y|G_k) \prod_{k=1}^c \prod_i P_k(\mathbf{x}_i|y_i) \quad (7)$$

where we ignore all the cross-community probability in order to simplify our model and reduce the cross-community influence.

In order to instantiate the probability of $P_k(Y|G_k)$ and $P_k(\mathbf{x}_i|y_i)$, we model them into a Markov random field, and according to the Hammersley-Clifford theorem [6], the probabilities can be instantiated as:

$$P_k(\mathbf{x}_i|y_i) = \frac{1}{Z_1} \exp\{\alpha_k \cdot \mathbf{f}_k(\mathbf{x}_i, y_i)\} \quad (8)$$

where α_k is the weight of feature vector function and $\mathbf{f}_k(\mathbf{x}_i, y_i)$ is the set of feature functions in the k^{th} community. Similarly,

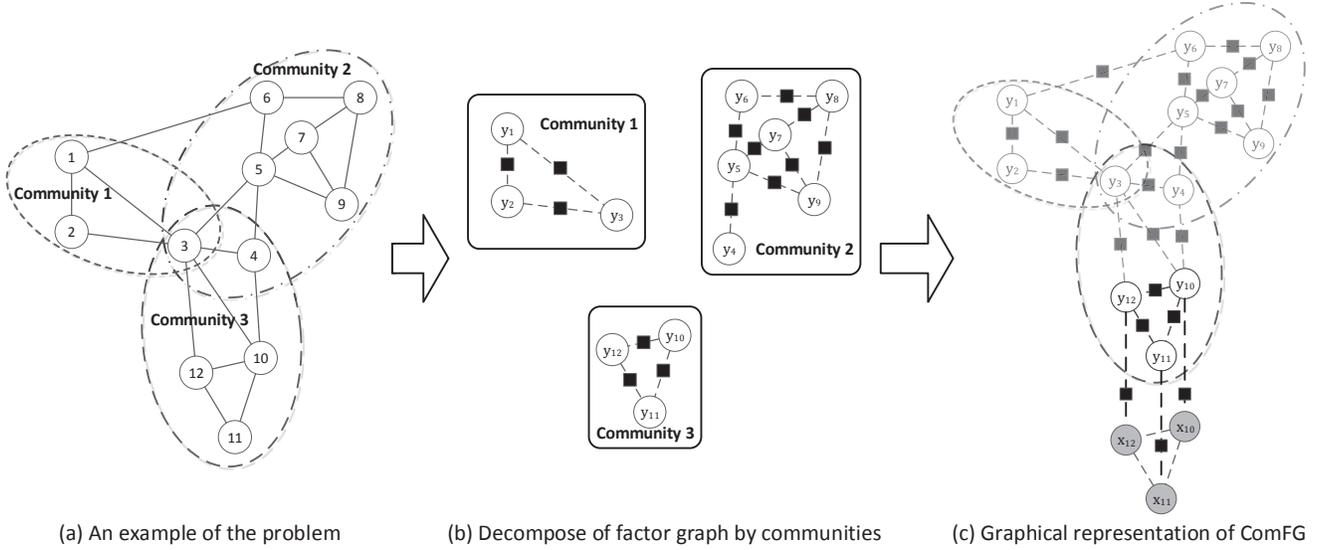


Figure 2: An example of the problem and graphical representation of ComFG model.

$$P_k(Y|G_k) = \frac{1}{Z_2} \exp\{\mu_k \cdot \mathbf{h}_k(y_i, y_j)\} \quad (9)$$

where μ_k is the weight of correlation vector function and $\mu_k \cdot \mathbf{h}_k(y_i, y_j)$ is the set of correlation functions in the k_{th} community.

Model Learning. We now learn to ComFG model by estimating the parameter configuration $\theta^* = (\{\alpha^*\}, \{\mu^*\})$ such that $\theta^* = \arg \max_{\theta} P(Y|\mathbf{X}, G)$ to optimize the log-likelihood of observed data. We can define the log-likelihood objective function as

$$\mathcal{O}(\theta) = \log \left(\prod_{k=1}^c P_{k\theta}(Y|G_k) \prod_{k=1}^c \prod_i P_{k\theta}(\mathbf{x}_i|y_i) \right) \quad (10)$$

A gradient method (Newton-Raphson method) is used to solved the objective function (Eq. 10). The gradient of each α_{kj} is

$$\frac{\delta \mathcal{O}(\theta)}{\delta \alpha_{kj}} = \mathbb{E}[f_{kj}(x_{ij}, y_i)] - \mathbb{E}_{P_{\alpha_{kj}}(x_{ij}, y_i)}[f_{kj}(x_{ij}, y_i)] \quad (11)$$

where $\mathbb{E}[f_{kj}(x_{ij}, y_i)]$ is the expectation of factor function $f_{kj}(x_{ij}, y_i)$ given the data distribution; and $\mathbb{E}_{P_{\alpha_{kj}}(x_{ij}, y_i)}[f_{kj}(x_{ij}, y_i)]$ is the expectation of factor function $f_{kj}(x_{ij}, y_i)$ under the distribution $P_{\alpha_{kj}}(Y|\mathbf{x}, G_k)$ given by the estimated model.

Similarly, for each μ_k is

$$\frac{\delta \mathcal{O}(\theta)}{\delta \mu_{kj}} = \mathbb{E}[h_{kj}(Y_c)] - \mathbb{E}_{P_{\mu_{kj}}(Y_c|X, G_k)}[h_{kj}(Y_c)] \quad (12)$$

where Y_c is short form of edge (y_s, y_t) ; $\mathbb{E}[h_{kj}(Y_c)]$ is the expectation of factor function $h_{kj}(Y_c)$ given the data distribution; and $\mathbb{E}_{P_{\mu_{kj}}(Y_c|X, G_k)}[h_{kj}(Y_c)]$ is the expectation of factor function $h_{kj}(Y_c)$ under the distribution $P_{\mu_{kj}}(Y_c|X, G_k)$ given by the estimated model.

We use Loopy Belief Propagation (LBP) to approximate the gradient and update θ iteratively and calculate the marginal distribution $P_{\mu_{kj}}(Y_c|X, G_k)$.

3.2 Feature Extraction

To train the proposed ComFG model, besides the homophily patterns, we also extract image attributes from two aspects. The first aspect is color features including saturation, brightness, color ration and dominant colors. Another aspect is social correlations which

Name	Description
Saturation	The mean and deviation of saturation
Brightness	The mean and deviation of brightness
Color Ratio	Cool color ratio and clear color ratio
Dominant Colors	The index of 81 HSV bins, where the HSV space is divided by 9 partitions along the Hue dimension, and three equal partitions along the other two dimensions
Uploading Time	The standard Unix time stamp is used to sorted images in time sequence
Community ID	The groups that the image are included in

Figure 3: Summary of all features

contain uploading time and community ID. The detail description lists listed in Figure 3.

4. EXPERIMENT RESULTS AND DISCUSSIONS

In this section, we first describe our experimental setup, then present the performance for several baselines and our model. Finally, we give some analysis and discussions.

4.1 Experimental Setup

Data set. The data set we used in this paper is collected from Flickr¹. The data set includes 150,000 images from 46972 groups and 2077 users. One user could belong to different groups and so is one image. In this work, we divide the data set into two parts evenly: both the training set and testing set include 75,000 images respectively.

Comparison Methods. We evaluate our proposed method with the following methods on Flickr data set.

- **SVM:** it uses the images' color features and homophily patterns to train a classification model and then predicts images' affects in the testing data. We use SVM-multiclass².

¹<http://www.flickr.com/>, Flickr is a photo sharing network

²<http://lightsvm.joachims.org/>

Data	Algorithm	Precision	Recall	F1-measure	Accuracy
Flickr	SVM	0.3759	0.3233	0.3476	0.3990
	LRC	0.3310	0.3470	0.3388	0.3720
	CRF	0.4203	0.3500	0.3819	0.4050
	FGM	0.4514	0.3710	0.4072	0.4210
	ComFG	0.4820	0.3740	0.4211	0.4440

Figure 4: Performance of different methods on Flickr data set

- **LRC:** it uses the same features as SVM to train a Logistic Regression Classification model and then predict images' affects in the testing data.
- **CRF:** it uses the same features as SVM to train a Conditional Random Field model to predict images' affects in the testing data.
- **FGM:** it trains a Factor Graph Model with partially labeled information with all image attributes (color features and social correlations) and homophily patterns.
- **ComFG:** the proposed model, which trains a Community-aware Factor Graph Model with all image attributes (color features and social correlations) and homophily patterns.

Evaluation Metrics. To quantitatively evaluate the performance of proposed method, we consider the following performance metrics:

- **Prediction accuracy.** We apply the learning model for predicting images' affects and evaluate its performance in terms of *Precision*, *Recall*, *F1-Measure*. We further analyze the factor contributions respectively.
- **Effects of homophily.** We use compare method to present the effect of different homophily patterns.
- **Scalability performance.** We evaluate the prediction performance by using different number of labeled data as the efficiency metrics.
- **Qualitative case study.** We use several case studies to further demonstrate the effectiveness of the proposed model.

All codes are implemented in C++ and JAVA, and all the evaluations are performed on an x64 machine with E7520 1.87GHz Intel Xeon CPU and 128GB RAM. The operation system is Microsoft Windows Server 2008 R2 Enterprise.

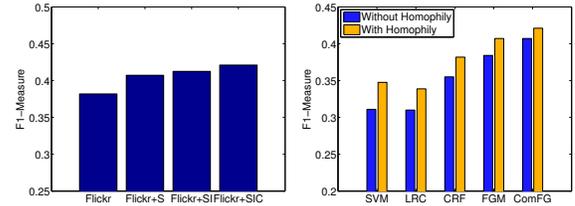
4.2 Performance Analysis

Prediction accuracy. Figure 4 lists the affects prediction results of the different methods on Flickr data set. The proposed method (ComFG) shows clearly the best performance than other methods. In terms of F1-Measure, ComFG achieves 7.35% improvement compared with SVM, 8.23% improvement compared with LRC, and 3.92% improvement compared with CRF. It demonstrates that besides color features, social correlations are also very helpful. Although FGM and ComFG can both incorporate color features and social correlations, as Flickr is a community-structured social network, the proposed ComFG achieves 3.06% improvement in Precision and 0.3% improvement in Recall compared to FGM (Figure 5). These results confirms the effectiveness of the proposed ComFG model.

Factor contribution analysis. To investigate the factor contributions in our problem, we first use only color features in ComFG and evaluate the performance by adding each of the other factors step by step into the model. Therefore, we can measure their contributions by the improvement of F1-Measure they achieved, as shown in Fig. 6(a). Besides color features, we have 3 types of factors: social correlations (S), individual homophily (I), and community

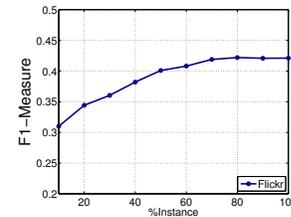
Affects	Precision		Recall		F1-measure	
	FGM	ComFG	FGM	ComFG	FGM	ComFG
Happy	0.4080	0.4158	0.5224	0.5321	0.4582	0.4668
Sad	0.5551	0.4810	0.4737	0.5050	0.5112	0.4927
Anger	0.3123	0.3275	0.1584	0.1870	0.2102	0.2381
Disgust	0.3586	0.3980	0.1324	0.1685	0.1934	0.2368
Surprise	0.5072	0.5015	0.5543	0.5663	0.5297	0.5319
Fear	0.3428	0.3610	0.3773	0.3671	0.3592	0.3640

Figure 5: Performance of FGM and ComFG on different affects



(a) Factor Contributions

(b) Effects of Homophily



(c) Scalability Performance

Figure 6: (a) Factor contribution analysis. Flickr stands for training a ComFG model with only the color features. Flickr+S stands for ComFG with social correlations. Flickr+SI stands for further involving individual homophily. ComFG+SIC stands for considering community homophily. (b) F1-Measure of different methods by considering the effect of homophily in Flickr network. (c) Performance by using different number of labeled data.

homophily (C). We can see different factors achieve different contributions. For example, in term of the F1-Measure, we get 3.5% improvement by adding social correlations, 0.3% more improvement by further adding individual homophily, and another 1.1% improvement by finally adding community homophily. These results confirm the effectiveness of the proposed factors.

Effects of homophily. In order to further show the effects of homophily, we use a compare experiment to present the power of homophily. Figure 6(b) shows the prediction accuracy in terms of F1-Measure on the different methods. It can be clearly seen that by incorporating the homophily factors, the prediction performance can be significantly improved 1.4%-3.4% for all the predictive methods. These results indicate the importance of modeling homophily in inferring affects from images in social networks.

Scalability performance. The accuracy results also depend on the size of training set for the initialization. A small number might result in high precision but low recall, while a large number might mean higher recall but would hurt the precision. Figure 6(c) shows how the average performance changes by varying the size of the training set. When the size of the training set is larger than 60,000, F1-Measure grows slowly, which indicates the rationality of using 75,000 images as training set in our experiments.

4.3 Qualitative Case Study

Now we give an interesting case study to further demonstrate the effectiveness of the proposed model. “Analogy Photography” and “Canon DSLR Users” are both of the most active photographers’ communities in Flickr, but they have different preferences. “Analogy Photography” advocate none-digital photos, while “Canon DSLR Users” are the enthusiasts of digital photos. Will the different community culture lead to different views on affects conveyed by images? Taking the images from the two communities as examples, we use the proposed model to infer affects of them. Some results are shown in Figure 7. Let’s focus on the “happy” and “sad” affects. For “happy” affect, we find its corresponding images from “Analogy Photography” mostly have brighter and warmer colors than those from “Canon DSLR Users”, and the blue channel of the color histograms is mostly larger than that in “Canon DSLR Users”. While for “sad” affect, the images contained in “Canon DSLR Users” are mostly with darker colors and higher contrast; and the red channel of the color histograms is mostly larger than that in “Analogy Photography”. So it can be clearly seen that people in different communities may have different opinions on colors and affects, but people in the same community have much more consistency. This case further indicates the rationality and importance of modeling homophily to propose the community-aware factor graph model for our task.

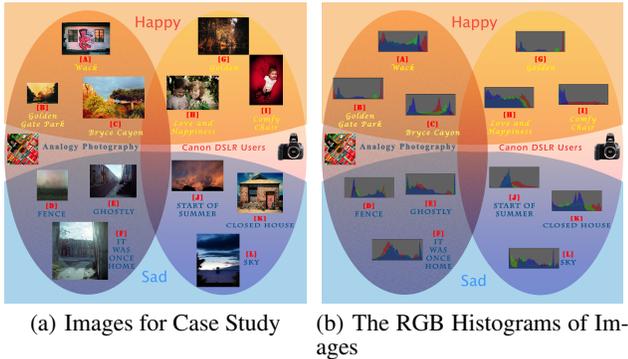


Figure 7: Case Study. Images from A-F comes from the “Analogy Photography” community in Flickr, and images from G-L are contained in the “Canon DSLR Users” community. The images and their corresponding color histograms are shown on the left and right side respectively. Contours with different shadows present the certain community or inferred affect.

5. RELATED WORK

Colors and Affects. Colors as described in [2], is the fundamental image attributes that reflect and impact the human affects. For the correlations between colors and affects, [14] studied the relationship between affective reactions and color hue, saturation and brightness using Pleasure-Arousal-Dominance affective model. [8] examined color-affect associations and found that affective responses to colors change with value and saturation. [10] made a summary about color features inspired by psychology and art theory, and further used the color features for affective images classification.

Social Attributes and Affects. Social attributes using in affective image classification always include ages, genders[3], cultures[3], races etc. [13] found that within 3 age groups, participants from the same group preferred consistent choices for color-affect pair comparison tasks. [11] revealed that there are correlations between color preference and interests of people. Although the importance of social attributes has been emphasized, most work on inferring

affects of images still only use visual attributes such as colors, composition, texture, and contents[10, 4], due to the lack of proper model to incorporate both the visual attributes and social attributes. In the next section, we will further give the problem definition before we introduce our ComFG model in details.

6. CONCLUSIONS

In this paper, we formally define patterns of homophily in individual and community levels, and then propose a community-aware factor graph model to infer affects from social images by leveraging the homophily patterns. The improved performance and case study reveal that the homophily plays an important role in effecting image affects in social networks.

Understanding the underlying homophily phenomenon can benefit many aspects, such as image retrieval by affective semantics and community-aware recommendation. So far, we have only studied homophily in interest-based online communities in this paper. As for the future work, it would be intriguing to find whether there is homophily behavior for communities based on other social factors, such as location and cultural background, to better capture the dynamic social ties.

7. REFERENCES

- [1] S. Aral and D. Walker. Identifying influential and susceptible members of social networks. *Science*, 337(6092):337–341, 2012.
- [2] F. Birren. *Color psychology and color therapy: A factual study of the influence of color on human life*. Kessinger Publishing, 2006.
- [3] A. Choungourian. Color preferences and cultural variation. *Perceptual and motor skills*, 26(3c):1203–1206, 1968.
- [4] R. Datta, D. Joshi, J. Li, and J. Z. Wang. Studying aesthetics in photographic images using a computational approach. In *Computer Vision–ECCV 2006*, pages 288–301. Springer, 2006.
- [5] P. Ekman, W. V. Friesen, M. O’Sullivan, A. Chan, I. Diacyoyanni-Tarlatzis, K. Heider, R. Krause, W. A. LeCompte, T. Pitcairn, P. E. Ricci-Bitti, et al. Universals and cultural differences in the judgments of facial expressions of emotion. *Journal of personality and social psychology*, 53(4):712, 1987.
- [6] J. M. Hammersley and P. Clifford. Markov field on finite graphs and lattices. *Unpublished manuscript*, 1971.
- [7] J. Jia, S. Wu, X. Wang, P. Hu, L. Cai, and J. Tang. Can we understand van gogh’s mood?: learning to infer affects from images in social networks. In *Proceedings of the 20th ACM international conference on Multimedia*, pages 857–860. ACM, 2012.
- [8] S. Kobayashi. The aim and method of the color image scale. *Color Research & Application*, 6(2):93–107, 1981.
- [9] G. Kossinets and D. J. Watts. Origins of homophily in an evolving social network. *American Journal of Sociology*, 115(2):405–450, 2009.
- [10] J. Machajdik and A. Hanbury. Affective image classification using features inspired by psychology and art theory. In *Proceedings of the international conference on Multimedia*, pages 83–92. ACM, 2010.
- [11] S. E. Palmer and K. B. Schloss. An ecological valence theory of human color preference. *Proceedings of the National Academy of Sciences*, 107(19):8877–8882, 2010.
- [12] J. Tang, S. Wu, and J. Sun. Confluence: Conformity influence in large social networks. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2013.
- [13] M. M. Terwogt and J. B. Hoeksma. Colors and emotions: Preferences and combinations. *The Journal of general psychology*, 122(1):5–17, 1995.
- [14] P. Valdez and A. Mehrabian. Effects of color on emotions. *Journal of Experimental Psychology: General*, 123(4):394, 1994.
- [15] X. Wang, J. Jia, P. Hu, S. Wu, J. Tang, and L. Cai. Understanding the emotional impact of images. In *Proceedings of the 20th ACM international conference on Multimedia*, pages 1369–1370. ACM, 2012.