Propagation of Facial Identities in a Social Network

Tao Wu
University of Maryland
College Park, MD 20742 USA
taowu@umiacs.umd.edu

P. Jonathon Phillips
National Institute of Standards and Technology
Gaithersburg, MD 20899 USA
jonathon.phillips@nist.gov

Rama Chellappa
University of Maryland
College Park, MD 20742 USA
rama@umiacs.umd.edu

Abstract

We address the problem of automated face recognition on a social network using a loopy belief propagation framework. The proposed approach propagates the identities of faces in photos across social graphs. We characterize performance in terms of structural properties of a social network. We propose a distance metric defined using face recognition results for detecting hidden connections. The result demonstrates that the constraints imposed by the social network have the potential to improve the performance of face recognition methods. The result also shows it is possible to discover hidden connections in a social network based on face recognition.

1. Introduction

A social network reflects the relationship structure among entities. It typically consists of different kinds of information, such as text, images and videos. The structure of a social network has been shown to play an important role in many fields such as marketing and epidemiology. It is an open question in face recognition, and computer vision in general, how algorithms can be adapted to solve vision problems on a social network. In this paper we address this question for face recognition algorithms.

Millions of facial images are uploaded to social network websites. The faces in these images are usually taken with point and shoot cameras or cell phones in unconstrained environments. This class of images is among the most challenging for face recognition. In this paper, we study how the structure of a social network can be used to improve the performance of automated face recognition algorithms.

We introduce a loopy belief propagation framework for face recognition to take advantage of the structure of social networks. The results of our experiments show that incorporating the structure of a social network does improve face recognition performance. We characterize algorithm performance as a function of the parameters of the social network. The parameters include the size of the network, number of friends/connections for each person, percentage of faces labelled, and face labelling errors in the network. We also examine the ability to detect hidden connections.

Measuring the performance in terms of a network's structure is vital to understanding the impact of face recognition, and computer vision in general, on social networks. To the best of our knowledge, this paper presents the first effort to characterize face recognition in terms of the structure of a social network.

Stone, et. al., [24, 25] showed that massive social media data can improve the performance of traditional computer vision and pattern recognition methods. Yu, et. al., [29] showed that hidden social connections could be found by analyzing the results of computer vision algorithms on social media data. A brief summary of related works is presented in Section 2.

In this paper, our primary contributions focus on characterizing the properties of the structure of a social network in improving face recognition performance. The proposed loopy belief propagation framework formulates the problem of face recognition on social networks as propagating identities of face images on social graphs. We also propose a distance metric which is defined using face recognition results for detecting hidden connections. The performance of the proposed method is analyzed in terms of graph structure, scalability, degrees of nodes, ability to correct labeling errors and discovering hidden connections.

2. Related works

2.1. The effects of social networks on computer vision

The social network structure and social media have wide applications in many fields. The social media contains not only face images, but also text and manual annotations. The information provided in various forms could be combined with computer vision methods. Recent efforts have shown the performance of face recognition and computer vision algorithms can be improved by incorporating available meta data into algorithms.

Stone, et al., [24, 25] automatically tagged facial images on a social network by combining face recognition results using a conditional random field model with social context, such as timestamp, geotag and other manual annotations. Dantone, et al., [7] extended this method to a practical automatic face recognition system for mobile devices. Choi, et al., [5] proposed a collaborative face recognition framework based on recursive polynomial models and sharing supervised identity information from users' feedbacks on face recognition for social network platforms. Another collaborative face recognition method for social network is presented in [4]. They proposed to fuse the results of a set of face recognition engines. Each classifier in the set of engines works on different features extracted from facial images. This work is further extended in [3], where the final face identities are obtained by fusing the results of the face recognition engine selected with the help of social context. Mavridis, et al., [17] described several algorithms that enhance face recognition by selecting multiple classifiers based on co-occurrence of faces in photos. Tseng, et al., [28] designed a photo identity suggestion method which relies only on the co-occurrence contexts, such as which user is manually tagged or left comments in albums. Poppe [22, 23] introduced several face labeling strategies based on the co-occurrence of faces in the same photos to infer their identities and validated the scalability of the strategies on large social networks. The social context information retrieved from communication, calender and collaborative applications and etc. are also able to improve the accuracy of face recognition on user-uploaded facial images [8]. These works show that traditional computer vision and pattern recognition algorithms could be improved using social networks.

2.2. The effects of computer vision on social networks

It is worth noting that the information obtained from computer vision algorithms enables improved social network analysis. Yu, et al., [29] presented a graph-cut based algorithm to discover hidden social connections using the results of face recognition and person tracking in images and videos captured by a pan-tilt-zoom (PTZ) camera system; Mavridis, et al., [17] proposed to predict friendship between people by counting the co-occurrence of faces in photos; Ding, et al., [9] illustrated an algorithm based on support vector regression using visual and auditory features and an affinity learning procedure to build the social network of characters in movies; another visual concept-based algorithm for discovering the social network of the characters in movies is reported in [10]; Minder, et al., [18] described a method for user re-identification in different social network sites based on the results of face recognition and text-attribute comparison. These works show that it is possible to discover new knowledge in a social network based on the results of computer vision algorithms.

2.3. Models and algorithms for social networks

In this paper, we focus on the role of social connectivity in propagating facial identities. We assume each person has a photo album consisting of facial images of him/herself and his/her friends. A face classifier can be trained for each person and tested on his/her friends. There are interactions between the classifiers on different subjects since the image labels in one's album could be updated using the outputs of classifiers trained on others' albums. We formulate the problem of identifying the unlabeled facial images based on the already labeled facial images and the connections in the social network as a belief propagation problem on an undirected graph. This is a popular method for graph-based inference.

There is a large body of research work on belief propagation and graphs. Belief propagation is an algorithm for inference on graphical models, such as Bayesian network and Markov random fields. Several exact and approximate Bayesian networks inference algorithms are summarized in [13]. Loopy belief propagation is an important method for Bayesian networks with loops. Murphy, et al., [20] empirically studied the application of belief propagation algorithms in networks with loops and suggested that good approximation could be obtained when the algorithm converges and the momentum could be helpful in reducing the oscillations. Acemoglu, et al., [2] employed Bayesian learning in social networks and studied the condition of asymptotic learning. They showed that expanding observations and unbounded private beliefs are sufficient conditions for asymptotic learning. Mossel, et al., [19] proposed a Bayesian model for iterative learning on social networks. They assumed that in a connected network, each agent estimates the status of variables by iteratively taking the optimal action given its belief and its neighbors' actions. Zheleva, et al., [30] explored the application of Markov random fields for inferring hidden attributes in social and affiliation networks. Everitt [12] applied a particle Markov chain Monte Carlo method to the problem of estimating the parameter of exponential random graphs from social network data. Tang, et al., [27] utilized loopy belief propagation to infer the type of social relationships for publication, email and mobile networks. These works show the effectiveness of Bayesian network and Markov random field methods in graphs.

In this paper, we formulate the face recognition problem on social networks as a belief propagation problem involving facial identities and social structures without any other social context information, and study the relationship between the structure of social networks and face recognition performance.

3. Propagation of facial Identities

3.1. Representation of a social network

A social network has a graph structure. In the graph, we model each person as a node, and connections between two people as an edge. In a social network, each person has connections to a set of friends. We consider the relationship between two people to be symmetric, and model the social network as an undirected graph (V, E), where V is the set of nodes that represent people, and E is the set of edges that represent the friendship between a pair of people. The degree of a node is the number of edges adjacent to it. Of interest to our investigation, are social networks where a person can upload photos. Since we are performing face recognition, we will assume that the uploaded photos contain only faces. In our model, we assume that each person uploads a set of face images that contain faces. We attach an album A_i to each node v_i that contain the set of face images. The album A_i consists of a set of face images: $A_i = \{I_{i,l}\}$, where $1 \le l \le n_i$ and n_i is the number of images in A_i.

3.2. Belief propagation

In our model, a face image could be initially labeled by the uploader or his/her friends. The probability of an image of person v_i appearing in album A_j depends on the distance between v_i and v_j on the graph. We are interested in estimating the identities of the unlabeled faces. There usually is a correlation between the images uploaded by friends, so we model the social network as a pair-wise Markov random field (MRF). The identities of the unlabeled images are represented as random variables in the MRF. The problem of inferring the identities of unlabeled images given the labeled images and the social network structure can then be formulated as a loopy belief propagation (BP) framework. Numerous papers have empirically demonstrated the performance of loopy BP algorithms [13, 14]; however, it has not been proven to converge on loopy networks.

We denote Al_i as an n_i-dimension random vector in order to represent the identities of all the images in album A_i,

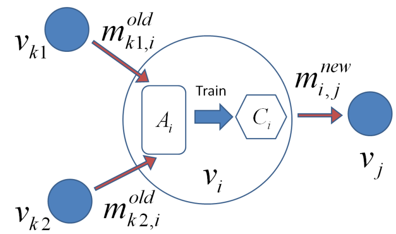

Figure 1. Illustration of message propagation on the network with albums of facial images.

such that:

$$Al_i = [L_{i,1}, \cdots, L_{i,l}, \cdots, L_{i,n_i}] \tag{1}$$

where $L_{i,l}$ is the identity of the image $I_{i,l}$.

A face recognition classifier C_i is trained for each node v_i with images in A_i and their corresponding labels Al_i, and is expected to give a probability distribution $C_i(L_{j,k})$ of the candidate identities for each image $I_{j,k}$.

Assuming that the labels $L_{i,l}$ are independent random variables, the joint probability distribution of the identities of all the images in the social network is given by:

$$P(Al_1, \cdots, Al_N) = \frac{1}{Z} \prod_{<i,j>} C_i(Al_j) \tag{2}$$

where Z is a normalizing factor and

$$C_i(Al_j) = \prod_{l=1}^{n_j} C_i(L_{j,l}) \tag{3}$$

Applying the max-product BP algorithm framework, the updated message $m^{new}_{i,j}$ from v_i to v_j is given by:

$$m^{new}_{i,j}(Al_j) = \max_{Al_i} C_i(Al_j) C_i(Al_i) \prod_{k \in Nbd(i)/j} m^{old}_{k,i}(Al_i) \tag{4}$$

where $Nbd(i)$ is the set of nodes that have direct connections to v_i, and $m_{i,j}(Al_j) = \prod_{l=1}^{n_j} C_i(L_{j,l})$. The message propagation is illustrated in Fig. 1.

The belief of the image $L_{i,l}$ in the album of v_i can be read out from the graph by:

$$b_i(L_{i,l}) \propto C_i(L_{i,l}) \prod_{k \in Nbd(i)/j} m_{k,i}(L_{i,l}) \tag{5}$$

If there are p possible identities for a facial image, the number of possible states of A_i is n_i^p. It becomes infeasible to compute over such a large number of possible states. We adopt a method which is similar to the particle filter method developed for BP [26] that resamples particles with high importance. The new message in (4) is updated using the identities with high probabilities. The images with high belief identities in the album of v_i will be added to the training set of C_i in the next iteration.

3.3. Classifiers

The proposed framework assumes that face recognition algorithms can be trained with a set of training images with their labels, and then output a probability distribution of candidate labels for each test images. Any classification method that satisfies this requirement can be adopted. Note that the proposed method can also be applied to other pattern recognition problems other than face recognition by using appropriate training and testing features. In our experiments, we use a Bayesian classifier [11] as a basis for our face recognition method.

For each node, a classifier is trained first using the initially labeled facial images in the node's album. Then the classifier is tested on all the images in albums associated with the node and its neighbor, and is re-trained in the next iteration using images with beliefs higher than a threshold in its album. These steps are repeated according to the loopy belief propagation algorithm so that the identities of facial images can be propagated through the social network.

3.4. Discovery of hidden connections

One of the interesting application is discovering hidden connections among people in a social network. A hidden connection means that people v_i and v_j are actually friends, but this connection is not explicitly shown in the graph structure. Discovery of hidden connections has many applications, such as modeling recommendations from friends in a social network website.

In order to determine if a relationship exists, we need a measurement of the relationship between two people, given some labels of facial images in their albums. A straightforward method is to examine the overlap between the given labels between two albums, for example, the percentage of the labels that appear in both albums. This method could be easily affected by wrong or ambiguous labels. In addition, as this method does not analyze the content of the images, it cannot make use of the unlabeled facial images.

We introduce an algorithm which measures the relationship based on face recognition results. The idea is to measure how well one person knows the facial images present in the album of the other. Consider a node v_i with an album A_i, and the beliefs $b_i(L_{i,l})$ of the facial images which are computed from (5). The probability of the candidate labels of images $P_i(L_{i,l})$ can be obtained by normalizing $b_i(L_{i,l})$ such that

$$P_i(L_{i,l} = id_k^{(i)}) = \frac{b_i(L_{i,l} = id_k^{(i)})}{\sum_k b_i(L_{i,l} = id_k^{(i)})} \tag{6}$$

where $id_k^{(i)} \in ID^{(i)} = \{id_1^{(i)}, \cdots, id_{n_i}^{(i)}\}$ is the set of possible candidate labels of image $I_{i,l}$.

A classifier C_j trained with A_j is tested on images in A_i. It yields a probability distribution $P_j(L_{i,l})$ of the candidate labels of all the images in A_i. Similarly, $P_j(L_{i,l})$ can be computed by normalizing $C_j(L_{i,l})$ as

$$P_j(L_{i,l} = id_k^{(j)}) = \frac{C_j(L_{i,l} = id_k^{(j)})}{\sum_k C_j(L_{i,l} = id_k^{(j)})} \tag{7}$$

where $id_k^{(j)} \in ID^{(j)} = \{id_1^{(j)}, \cdots, id_{n_j}^{(j)}\}$ is the set of possible candidate labels given by C_j.

The distance between v_i and v_j is measured by comparing these two probability distributions. We use the Kullback-Leibler divergence as a measurement of this distance. It is possible that there is no common support between $P_i(L_{i,l})$ and $P_j(L_{j,l})$, in which case the Kullback-Leibler divergence is not defined. Hence we apply Laplacian smoothing [16] such that

$$P_i(L_{i,l} = id_k) = \frac{b_i(L_{i,l} = id_k) + \alpha}{\sum_k b_i(L_{i,l} = id_k) + \alpha d} \tag{8}$$

$$P_j(L_{i,l} = id_k) = \frac{C_j(L_{i,l} = id_k) + \alpha}{\sum_k C_j(L_{i,l} = id_k) + \alpha d} \tag{9}$$

where $id_k \in ID^{(i)} \cup ID^{(j)}$, $d = |ID^{(i)} \cup ID^{(j)}|$, and α is the smoothing parameter. The KL divergence can then be computed as:

$$kld_{j,i} = \frac{1}{n_i} \sum_{l=1}^{n_i} \sum_k \ln\left(\frac{P_j(L_{i,l} = id_k)}{P_i(L_{i,l} = id_k)}\right) P_i(L_{i,l} = id_k) \tag{10}$$

and the score $s_{i,j}$ of the connection between v_i and v_j can be defined as:

$$s_{i,j} = \frac{kld_{i,j} + kld_{j,i}}{2} \tag{11}$$

We threshold the score to determine if a connection between two nodes exists.

4. Experimental results

In this section, we first discuss the dataset we used in our experiments. Then we present experimental results that characterize the performance of our method on graph structure networks, and discuss dependence on factors such as scalability, degrees of nodes, ability to correct labeling errors and discovery of hidden connections.

4.1. Dataset

To the best of our knowledge, there is no publicly available database that provides facial images as well as the social connections among them. Hence, we use a publicly available social network dataset and a facial image database to generate a social network of facial image dataset.

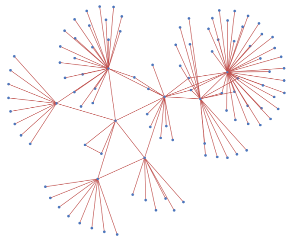

Figure 2. An example of a local structure of the social network extracted from the Stanford Large Network Dataset Collection [15].

The Stanford Large Network Dataset Collection [1] in Stanford Network Analysis Platform (SNAP) [1] is a publicly available dataset which provides network structures of large networks datasets, including the network data collected from real online social networks. It contains thousands of nodes and millions of edges. An example of the social network structure is shown in Fig. 2. We extracted a random subset of nodes and their edges to construct a synthetic social network.

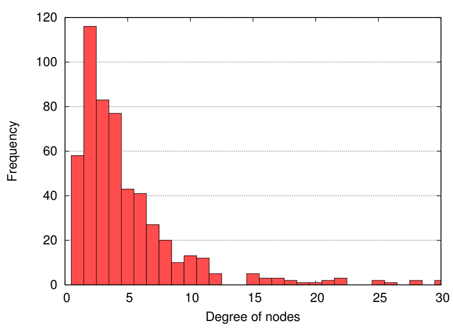

Figure 3. The distribution of the degrees of the nodes.

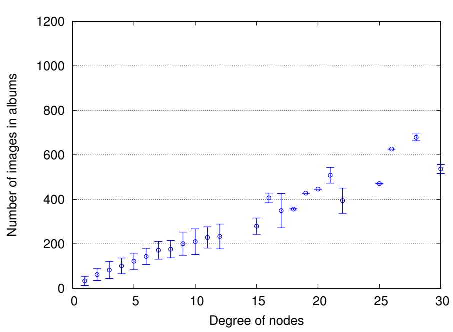

Figure 4. The means and standard deviations of the number of images in the albums of the nodes of different degrees.

We are interested in the impact of the connections in a social network on face recognition algorithms. Thus we use cropped facial images to avoid other interferences such as the uncertainty in the performance of face detection algorithms. The FRGC2.0 [21] is one of the largest facial image database, which consists of 39,328 facial images collected from 568 subjects. The facial regions are extracted from original facial images using the eye coordinates provided in the database. We use this face image dataset to generate the album for each node in the social network since it enables the result to reveal the effects of the structure of the network without interference from other factors such as poses and aging.

Each node is randomly assigned a unique identity of the subjects in FRGC. We observe that in online social networks, the images of one user often appear in the user or the user's friends' albums, and seldom appears in strangers' albums. So we generate the album for each node such that the probability of an image I_i of v_i that appears in the album A_j of v_j is a non-negative decreasing function of the distance between two nodes in the graph:

$$P(I_i \in A_j) = f(d(v_i, v_j)) \tag{12}$$

The exponential distribution function is a good candidate for the non-negative descreasing function. The distribution of the degrees of the nodes in the social graph is shown in Fig. 3. The means and standard deviations of the number of images in the albums of the nodes of different degrees are shown in Fig. 4. It shows that the size of the album of each node is proportional to the degree of the nodes since the images of all the friends of a node could appear in its album.

4.2. Results

4.2.1 Impact of social networks

We randomly select a subset of the images and assume they are labeled initially. The proposed method is then tested on the dataset. To measure the impact of the structure of social networks on performance, we compare rank 1 performance with two control network models. The control network models are:

S1) apply the face recognition algorithm individually on each album, i.e., the structure of the network is not exploited and no message is propagated across different albums. The results obtained in this situation are used as baseline 1.

S2) apply the face recognition algorithm on the union of the albums of all the nodes, i.e., all the albums are merged to a dataset such that the initially labeled images serve as the training set and other images serve as the testing set. The classifier does not use any information from the social network. The results obtained in this situation are used as baseline 2.

We empirically study the effect of the percentage of the initially labeled images on the performance of face recognition methods. At each percentage of initially labeled images, all the methods are tested on three randomly generated

[1]The identification of any commercial product of trade name does not imply endorsement of recommendation by NIST.

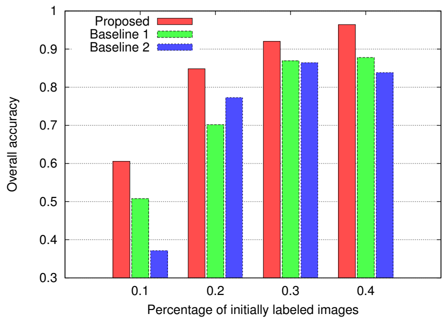

Figure 5. The overall rank-1 identification accuracies of the proposed methods and baselines S1 and S2. Performance is also characterized by the percentage of the initially labeled images.

datasets. The average overall rank-1 identification accuracies on unlabeled images using different methods with different percentage of initially labeled images are shown in Fig. 5.

The result demonstrates that compared to the traditional method that does not model the connectivity in a social network, the performance of face recognition is improved by exploring the structure of the social network. In situation S1, the Bayesian classifiers are applied locally to each album individually. The number of possible identities is fewer in a local album than the whole image set, which may be helpful to the local classifiers. In situation S2, the union of all the albums could attain more labeled training samples, which is helpful for training a single classifier. These may be the reasons that the performance of Bayesian classifiers are close under S1 and S2. However, when applied to large scale social networks with millions of subjects and training samples, both complexity and computational load will increase significantly. Thus face recognition under scenario S2 may be not feasible in real social networks.

4.2.2 Scalability over the size of the graph

Scalability is an important attribute for social network applications. We tested the performance of the proposed method on social graphs with different sizes.

We randomly extract a sub-network with 50 to 550 nodes from the SNAP dataset. Similarly, we generate the albums for the nodes. In this experiment, we assume that 40% of the images are initially labeled and test the proposed method and the baseline algorithm (scenario S2). The experiments are repeated 5 times for each set of nodes and the average overall rank-1 identification accuracies are shown in Fig. 6.

The result shows that the performance of the proposed method is stable as the size of the graph grows, while the performance of the traditional method slightly decreases. Since each node has its own album, adding new nodes to the graph introduces new identities and new facial images.

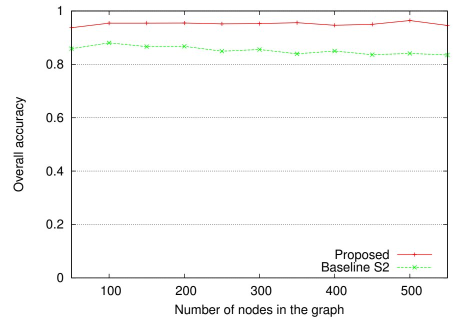

Figure 6. The overall rank-1 identification accuracy on graphs with different number of nodes of the proposed method and S2 baseline.

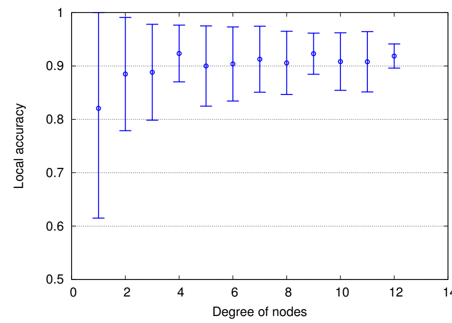

Figure 7. The means and standard deviations of the local (within albums) rank-1 identification accuracies at nodes of different degrees.

This increases both the classification difficulty and the required computing resources for a global classifier. Under the proposed framework, the classification task is always restricted to a local region which might be the reason that the performance is stable. This suggests that the proposed method is scalable to large scale social networks.

4.2.3 Effects of the degrees of the nodes

The degree is an important property for nodes. A node with a large degree implies that the node has many neighbors and that its album contains facial images from a large number of different subjects. We empirically study the relationship between the degrees of the nodes and the performance of face recognition.

The means and standard deviations of the local rank-1 accuracies on nodes with different degrees are shown in Fig. 7, assuming 30% of the images have been initially labeled. The result shows that as the degree of the node increases, the local performance increases and the standard deviation of the performance decreases. It also shows the performance tends to be stable with high accuracy on nodes with high degrees. This may be because such nodes have more facial images in their albums and the classifiers on these nodes can be better trained. Because there are too few nodes with degree greater than or equal to 15, we do not report their

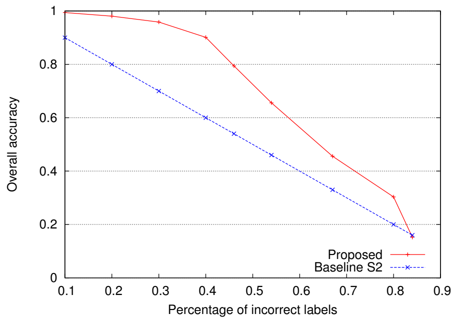

Figure 8. The overall rank-1 identification accuracy when different percentage of images have incorrect initial labels.

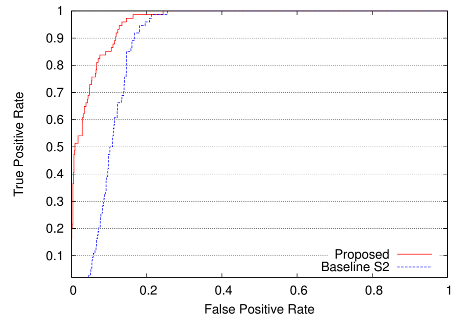

Figure 9. The ROC curves of detecting hidden connections between nodes.

performance.

4.2.4 Correction of incorrectly labeled images

For images on a social network website, the initial identities are usually manually labeled by users. It is possible that some labels could be incorrect. Determination of facial identity when there are incorrectly labeled faces in the training images is an open question. We are interested in how the proposed method performs when the training data contains errors. This problem is different from the partial label problem [6], since there is only one candidate label, either correct or incorrect for an image; while the partial label problem assumes that there are initially a set of candidate labels for an image and one of them is the correct label.

We first set the labels of all the images as the grondtruth. Next we randomly select a percentage of the images and change their labels. The proposed method and the traditional classification method are then tested on the datasets which contain errors. The overall rank-1 identification accuracies of different methods under different percentage of initially incorrect labels are shown in Fig. 8.

The result shows that when there are errors in the training set, the classifiers learns the errors, which makes the performance of the baseline method decline to a diagonal line (Fig. 8). In the proposed framework, the classification results of the same image from multiple classifiers that are trained on different training set are fused. This gives the algorithm the capability to correct some of the incorrect labels. It also shows that the rank-1 identification accuracy of the proposed method is greater than 90% when no more than 40% of the images are incorrectly labeled.

4.2.5 Discovery of hidden connections

In order to evaluate the ability for discovering the hidden connections, we first extract the sub-network and generate a set of albums. Next we randomly remove a number of connections from the graph so that we have a groundtruth. These connections are expected to be detected with high confidence, thus we may use ROC curves to characterize the performance. The proposed method is tested on the graph with removed edges. The score is computed between each pair of nodes when the connection is not known. The hidden connection can then be detected by selecting a threshold on the scores of the connections.

We take a straightforward method which examines the percentage of the labels that appear in both albums as a baseline. The ROC curves of detecting the hidden connections using the proposed method and the baseline are shown in Fig. 9. It shows that the proposed method is capable of detecting the hidden connections effectively. The detected result can be used in applications such as recommending friends on a social network website.

5. Conclusion

In this paper, we study the problem of face recognition across social network by formulating it as a belief propagation problem. We build multiple synthetic social network facial image dataset by combining a publicly available social network and facial image datasets. The results demonstrate that the structure of the social network contains information that improves the performance over traditional face recognition scenarios. As the degree of a node increases, the recognition accuracy on its album increases and the variance of the performance accuracy decreases. In addition, it is possible to discover hidden connections in the social network based on face recognition results. Although we adopt a Bayesian classifier as the face recognition technique in our experiments, other face recognition techniques can be integrated in the proposed framework.

One limitation of the proposed method that since the classifier are trained and tested locally on the social network, it is difficult to identify if a user uploaded facial images that belong to someone not close to him or her on the social network. A related extension is to move from a closed to open universe model for the faces on a network.

References

[1] Stanford large network dataset collection. *[Online]*, 2012. 5

[2] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar. Bayesian learning in social networks. Working Paper 14040, National Bureau of Economic Research, May 2008. 2

[3] J. Y. Choi, W. De Neve, K. Plataniotis, and Y. Ro. Collaborative face recognition for improved face annotation in personal photo collections shared on online social networks. *Multimedia, IEEE Transactions on*, 13(1):14 –28, feb. 2011. 2

[4] J. Y. Choi, W. De Neve, Y. M. Ro, and K. N. Plataniotis. Face annotation for personal photos using collaborative face recognition in online social networks. In *Proceedings of the 16th international conference on Digital Signal Processing*, DSP'09, pages 240–247, Piscataway, NJ, USA, 2009. IEEE Press. 2

[5] K. Choi, H. Byun, and K.-A. Toh. A collaborative face recognition framework on a social network platform. In *Automatic Face Gesture Recognition, 2008. FG '08. 8th IEEE International Conference on*, pages 1 –6, sept. 2008. 2

[6] T. Cour, B. Sapp, and B. Taskar. Learning from partial labels. *J. Mach. Learn. Res.*, 12:1501–1536, July 2011. 7

[7] M. Dantone, L. Bossard, T. Quack, and L. V. Gool. Augmented faces. In *IEEE International Workshop on Mobile Vision (ICCV 2011)*, 2011. 2

[8] H. A. David Petrou, Andrew Rabinovich. Facial recognition with social network aiding. *U.S. Patent 8121618*, 2011. 2

[9] L. Ding and A. Yilmaz. Learning relations among movie characters: a social network perspective. In *Proceedings of the 11th European conference on Computer vision: Part IV*, ECCV'10, pages 410–423, Berlin, Heidelberg, 2010. Springer-Verlag. 2

[10] L. Ding and A. Yilmaz. Inferring social relations from visual concepts. In *Proceedings of the 2011 International Conference on Computer Vision*, ICCV '11, pages 699–706, Washington, DC, USA, 2011. IEEE Computer Society. 2

[11] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification (2nd Edition)*. Wiley-Interscience, 2000. 4

[12] R. G. Everitt. Bayesian parameter estimation for latent markov random fields and social networks. *CoRR*, abs/1203.3725, 2012. 2

[13] H. Guo and W. Hsu. A survey of algorithms for real-time bayesian network inference. In *In In the joint AAAI-02/KDD-02/UAI-02 workshop on Real-Time Decision Support and Diagnosis Systems*, 2002. 2, 3

[14] T. X. Han, H. Ning, and T. S. Huang. Efficient nonparametric belief propagation with application to articulated body tracking. In *Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Volume 1*, CVPR '06, pages 214–221, Washington, DC, USA, 2006. IEEE Computer Society. 3

[15] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph evolution: Densification and shrinking diameters. *ACM Trans. Knowl. Discov. Data*, 1(1):2, 2007. 5

[16] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, New York, NY, USA, 2008. 4

[17] N. Mavridis, W. Kazmi, and P. Toulis. Friends with faces: How social networks can enhance face recognition and vice versa. In A. Abraham, A.-E. Hassanien, and V. Snel, editors, *Computational Social Network Analysis*, Computer Communications and Networks, pages 453–482. Springer London, 2010. 2

[18] P. Minder and A. Bernstein. Social network aggregation using face-recognition. In *ISWC 2011 Workshop: Social Data on the Web*, Bonn, Germany, 2011. 2

[19] E. Mossel and O. Tamuz. Efficient Bayesian learning in social networks with Gaussian estimators. *Computing Research Repository*, Feb. 2010. 2

[20] K. P. Murphy, Y. Weiss, and M. I. Jordan. Loopy belief propagation for approximate inference: An empirical study. In *In Proceedings of Uncertainty in AI*, pages 467–475, 1999. 2

[21] P. J. Phillips, P. J. Flynn, T. Scruggs, K. W. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 947–954, 2005. 5

[22] R. Poppe. Scalable face labeling in online social networks. In *FG*, pages 566–571. IEEE, 2011. 2

[23] R. Poppe. Facing scalability: Naming faces in an online social network. *Pattern Recogn.*, 45(6):2335–2347, June 2012. 2

[24] Z. Stone, T. Zickler, and T. Darrell. Autotagging facebook: Social network context improves photo annotation. In *Computer Vision and Pattern Recognition Workshops, 2008. CVPRW '08. IEEE Computer Society Conference on*, pages 1 –8, june 2008. 1, 2

[25] Z. Stone, T. Zickler, and T. Darrell. Toward large-scale face recognition using social network context. *Proceedings of the IEEE*, 98(8):1408 –1415, Aug. 2010. 1, 2

[26] E. B. Sudderth, A. T. Ihler, M. Isard, W. T. Freeman, and A. S. Willsky. Nonparametric belief propagation. *Commun. ACM*, 53(10):95–103, Oct. 2010. 3

[27] W. Tang, H. Zhuang, and J. Tang. Learning to infer social ties in large networks. In *Proceedings of the 2011 European conference on Machine learning and knowledge discovery in databases - Volume Part III*, ECML PKDD'11, pages 381–397, Berlin, Heidelberg, 2011. Springer-Verlag. 3

[28] C.-Y. Tseng and M.-S. Chen. Photo identity tag suggestion using only social network context on large-scale web services. In *Proceedings of the 2011 IEEE International Conference on Multimedia and Expo*, ICME '11, pages 1–4, Washington, DC, USA, 2011. IEEE Computer Society. 2

[29] T. Yu, S. N. Lim, K. Patwardhan, and N. Krahnstoever. Monitoring, recognizing and discovering social networks. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1462–1469, June 2009. 1, 2

[30] E. Zheleva, L. Getoor, and S. Sarawagi. Higher-order graphical models for classification in social and affiliation networks. In *NIPS Workshop on Networks Across Disciplines: Theory and Applications*, 2010. 2