Discriminative subgraphs for discovering family photos

Changmin Choi¹, YoonSeok Lee¹, and Sung-Eui Yoon¹ (✉)

© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract We propose to use discriminative subgraphs to discover family photos from group photos in an efficient and effective way. Group photos are represented as face graphs by identifying social contexts such as age, gender, and face position. The previous work utilized bag-of-word models and considered frequent subgraphs from all group photos as features for classification. This approach, however, produces numerous subgraphs, resulting in high dimensions. Furthermore, some of them are not discriminative. To solve these issues, we adopt a state-of-the-art, frequent subgraph mining method that removes non-discriminative subgraphs. We also use TF-IDF normalization, which is more suitable for the bag-of-word model. To validate our method, we experiment in two datasets. Our method shows consistently better performance, higher accuracy in lower feature dimensions, compared to the previous method. We also integrate our method with the recent Microsoft face recognition API and release it in a public website.

Keywords image classification; subgraph mining; social context; group photographs

1 Introduction

Recent studies on image classification focus on object and scene classification. They show remarkable performance thanks to the improvement of image features such as convolutional neural network (CNN) [1]. These image features are built from pixel-level descriptors, and may be not enough to describe group photos, since classifying group photos requires to utilize more semantic information like relations, events, or activities. Interestingly,

humans can classify types (e.g., friends and family) of group photos without much training, because we can estimate a variety of social contexts such as age, gender, proximity, and place, by observing face, position, clothing, and other objects.

Once we identify the social context on group photos, we can use this information for various applications. One application is to control privacy of shared images in various social websites (e.g., Facebook). People share images sometimes without much consideration on what information shared images can deliver to other people. When we identify that a shared group photo is a family photo containing children, we may wish to share that image to a small circle of persons, e.g., relatives, instead of publicly.

For classifying group photos, Chen et al. [2] proposed a method to categorize group photos into family and non-family types. This method assumes that annotations about age, gender, and face position are well-estimated beforehand by using existing face detection and statistical estimation derived from the pixel context. On top of that, they proposed to use a social-level feature named as Bag-of-Face-subGraph (BoFG) to represent group photos by graphs. For constructing BoFGs, a mining algorithm extracting frequent subgraphs is adopted. This is based on the assumption that prominent social subgroups captured in group photos can be identified by looking at frequently appearing subgraphs.

While the prior method enlightens an interesting research direction of classifying group photos, it has certain drawbacks. It first requires a user-specified threshold to determine the number of feature dimensions in a training phase. Furthermore, as we have more frequent subgraphs by having more feature dimensions, we also raise the probability as a

¹ Korea Advanced Institute of Science and Technology (KAIST), 291 Daehak-ro, Yuseong-gu, Daejeon, Republic of Korea. E-mail: sungui@gmail.com (✉).

Manuscript received: 2016-02-01; accepted: 2016-04-13

side effect that more non-discriminative subgraphs are selected due to repetitive and redundant patterns. In other words, thresholding the number of subgraphs with the frequency criterion alone can cause a scalability problem.

Main contributions. To overcome these issues, we survey the state-of-the-art subgraph mining techniques, and propose to use a subgraph mining technique, CORK, that identifies discriminative subgraphs and culls out redundant subgraph generations. We also propose to use a TF-IDF, a widely-used feature normalization for the bag-of-word models, to our BoFG feature.

To validate benefits of our method in terms of classifying family and non-family types of group photos, we have tested the prior and our methods in two different datasets (Fig. 1) including the public dataset [3]. Overall, our method shows higher accuracy with less dimensions over the prior method. Furthermore, our method does not require a manually tuned threshold for computing dimensions of our BoFG features.

We have also integrated our method with the face

Fig. 1 We test our method against a new, extended dataset consisting of (a) non-family and (b–e) different family types. Our method achieves the highest accuracy, 79.34%, with 90 dimensions, while the state-of-the-art method achieves 76.8% with 1000 dimensions.

API¹ of Microsoft Project Oxford and released it at our demo site². In this system (Fig. 2), users can test their own group images and see how well our method performs with them.

2 Related work and background

We review prior approaches that are related to our method.

2.1 Social context in photographs

Social contexts contain various information such as clothing, age, gender, absolute or relative position,

Fig. 2 Our demo site using the proposed classification method.

¹<https://www.projectoxford.ai/face/>.

²<http://is-fam.net/>.

face angle, gesture, body direction, and so on. They have been widely used to recognize people and groups [2, 4, 5]. Several works analyzed the contexts to study the structure of scenes in group photos [3, 6, 7]. Some researchers utilized them to classify group types [2, 4, 8, 9], retrieve similar group photos [10–12], discover social relations [5, 13], or predict occupations [14].

Pixel contexts in addition to the social contexts have been used together to recognize a type of group photos [4]. Some of well-known pixel-level features include SIFT [15], GIST [16], CNN [1], etc. Social-level features can be estimated by face detection, clothing segmentation, or partial body detection.

2.2 Frequent subgraph mining

Our work is based on identifying subgraphs from a graph representing the relationship between people shown in group photos. Frequently appearing subgraphs provide important cues on understanding graph structures and similarity between different graphs. As a result, mining frequent subgraphs has been widely studied [17]. For various classification, frequent subgraph mining has been used in training and test phases to build a social-level feature, as used in classifying family and non-family photo types [2].

We have found that extracted subgraphs significantly affect classification accuracy. There are two simple strategies to explore subgraphs in a database: (1) BFS-based and (2) DFS-based approaches [17]. The BFS-based algorithm has been less used recently due to its technical challenges in generating candidates and pruning false positives. More advanced techniques focus on efficient candidate generation, since the subgraph isomorphism test is an NP-complete [18]. Recent successful algorithms proceed based on depth-first search and pattern growth [17], i.e., subgraph growing. Our method is also based on the DFS-based strategy, and uses canonical labels to avoid the scalability issue. We additionally measure the discriminative power of each subgraph during the pattern growth.

2.3 Graph-based image editing

In this work, we use graphs and histograms of their subgraphs for discovering family photos. Interestingly, there have been many graph-based approaches for image extrapolation [19],

interpolation [20], image segmentation [21], representations [22], etc. While these applications are not directly related to our classification problem, utilizing histograms of subgraphs could be useful in these applications, e.g., better graph matching for extrapolation.

3 Backgrounds on social subgraphs

In this section, we give the background of using BoFG features for group photo classification.

Chen et al. [2] proposed BoFG features for group photo classification. This method constructs face graphs (Fig. 3) and uses their subgraphs to describe various social relationships. BoFG is analogous to the bag-of-words model of text retrieval. For example, a text corpus corresponds to a group photo album, a document to an image, and a word to a subgraph in a face graph, respectively. The main difference between these models is that the bag-of-words model performs clustering over all vectors in order to obtain a codebook, whereas BoFG performs frequent subgraph mining over all the face graphs.

(a) Generating a group photo to a face graph

Age range	0-2	3-7	8-12	13-19	20-36	37-65	66+
Female	1	5	10	16	28	51	75
Male	1	5	10	16	28	51	75

(b) 14 types of vertices

	0	1	2	3
Order distance	—	— + —	— + + —	— + + + —

(c) 4 types of edges

Fig. 3 (a) Representing an image as a face graph using (b) 14 vertex types and (c) 4 edge types.

Attributes of group members enable us to discriminate the type of groups, although we do not even know their names or relationships. In addition, understanding each one's position is informative to infer physical and relationship closeness among people. Chen et al. [2] showed that only knowing gender, age, and face positions as attributes of group members works effectively for a binary classification of family and non-family photos. Our approach is also based on this approach, and represents a group photo into a face graph, elaborated below.

Face graphs. Figure 3 illustrates an example of representing a group photo to a face graph. Each node of the graph corresponds to each person in the group photo, and is associated with a vertex label describing age and gender. Each edge between two nodes encodes relative position between two people.

There are 14 different types describing age and gender for each vertex label. The age ranges from 0-year-old to 75-year-old, and is categorized into seven age types. There are two gender types, male and female, and they are visualized by square and circle, respectively in Fig. 3(b). The combinations of age and gender result in 14 different types. Identifying faces and their attributes have been well studied [23, 24], and APIs of performing these operations are available, as mentioned in Section 1.

Most previous works used the Euclidean distance in image space, i.e., pixel distance, to measure the closeness between persons in group photos [3, 5, 12, 13]. Unfortunately, it has been known not to be invariant to scales of images, faces, distance to camera, or the orientation angle of a face. Instead, we use an order distance that indicates how close people stand with each other. The order distance has been demonstrated to be more stable over the pixel distance in terms of various factors [2]. The order distance is computed as the path length among vertices on a minimum spanning tree (MST) generated from a face graph. Such order distance is used for each edge label such as Fig. 3(c).

Bag-of-Face-subGraph (BoFG). Once we represent group photos into face graphs, we extract frequent subgraphs and regard them as BoFG features for classification. BoFG has been proposed to be a useful feature to compare structures of group photos. It helps to infer a type of a group by using substructures of groups. For example, in Fig. 3,

edges between two vertices of 28_f and 28_m (i.e., mother–father relationship), and between 28_f and 5_m (i.e., mother–son relationship) provide additional information on social relationship over each node of those edges; f and m represent female and male gender types, respectively.

Subgraph enumeration via gSpan. The prior work regarded frequent subgraphs as BoFG features, and generated such subgraphs by frequent subgraph mining, specifically, the gSpan method [25]. Most prior approaches of frequent subgraph mining [17] initially generate candidates of frequent subgraphs and adopt a pruning process to remove false positives. The pruning process, unfortunately, has a heavy computational cost, because it requires subgraph isomorphism testing.

gSpan adopted in the prior classification system [2] ameliorated this computational overhead issue by utilizing two techniques, DFS lexicographic order and minimal DFS code. Specifically, we first traverse an input graph, G , in a depth-first search (DFS) and assign an incrementally increasing visiting order to a newly visited vertex. Whenever we traverse an edge from v_m to v_n of the graph G , we represent the traversed edge into a 5-tuple DFS code:

$$G = (m, n, L_m, L_{(m,n)}, L_n) \quad (1)$$

where m and n are vertex indices computed by the visiting ordering during the DFS traversal, L_m and L_n are vertex labels of v_m and v_n , respectively, and $L_{(m,n)}$ is a edge label associated with the edge.

A graph, however, can have multiple DFS codes depending on traversal orders of vertices and edges. gSpan particularly allows the DFS lexicographic order computed from labels, $L_m, L_{(m,n)}, L_n$, of vertices and edges, and uses the DFS code corresponding to the minimal lexicographic order from the graph G . In this way, we can remove redundant subgraphs and maintain a subgraph among its isomorphic subgraphs.

To check the subgraph isomorphism, we simply look at the DFS code of a subgraph, G_s , to see whether the code is equal to or bigger than ones generated by prior subgraphs. If so, this indicates that G_s is a redundant subgraph, which is isomorphic to a prior subgraph. An illustration of generating DFS codes and pruning process is shown in Fig. 4.

To define frequently appearing subgraphs, gSpan requires a user defined parameter, known as

Fig. 4 This figure shows a process of generating all the subgraphs having one edge or more in gSpan. During the enumeration of subgraphs, gSpan prunes subgraphs once their DFS codes are equal to or bigger than prior ones. We highlight three subgraphs labelled (a), (b), and (c), and their DFS codes in below. 5_f is a 5-year-old female, while 5_m is a 5-year-old male. Let the lexicographic orders of vertex and edge be $5_f < 5_m < 28_f < 28_m$ and $0 < 1 < 2$. Note that the subgraphs (b) and (c) are isomorphic to each other. However, the subgraph of (c) is not a minimal DFS code because it is bigger than that of (b). In this manner, the search space can be pruned; the dotted subgraphs are pruned during the DFS-based expansion.

minimum frequency. We consider all different subgraphs whose frequency counts are bigger than the minimum frequency to be features of the BoFG.

The aforementioned method focuses on extracting frequency-based subgraphs and has some limitations for graph classification. Extracted frequent subgraphs in this approach may not show structural differences between classes. This is a similar problem even in the text classification. For instance, “a” and “the” are most commonly appearing words, but are not discriminative words for document classification. Moreover, the minimum frequency of subgraphs for defining BoFGs should be picked through a tedious trial-and-error approach for achieving high accuracy.

To address these drawbacks of using frequently appearing subgraphs, we propose to use discriminative subgraphs, adopt a recent subgraph mining method, CORK [26], extracting such discriminative subgraphs, and apply it to our classification problem of group photos. Additionally, we further improve the classification accuracy by adopting and tailoring the TF-IDF normalization scheme to our problem.

4 Our approach

In this section, we explain our approach for classifying group photos into family and non-family types.

4.1 Overview

Figure 5 shows the overview of our method. As offline process, we first generate face graphs from group photos in a training set and extract discriminative subgraphs as Bag-of-Face-subGraph (BoFG) features from face graphs. We utilize family and non-family labels associated with training images. We then extract a BoFG feature for each photo and normalize the feature by using the TF-IDF weighting. Through discriminative learning, we finally construct an SVM classifier.

When a query image is provided, we represent it to a face graph, and extract and normalize a BoFG feature from the graph. We then estimate a query’s label by utilizing the pre-trained classifier.

Our work adopts face graphs and their subgraphs as the BoFG features for the classification problem (Section 3). For achieving higher accuracy in an efficient manner, we additionally propose using discriminative subgraphs (Section 4.2) inspired by a recent near-optimal selection method [26]. We also normalize BoFG features using the term frequency and inverse document frequency, i.e., the TF-IDF weighting scheme (Section 4.3).

4.2 Discriminative subgraphs mining

We would like to identify discriminative subgraphs that are characteristic features in each category. We have identified similar issues from data mining, and found that CORK [26] works well for our problem.

CORK considers statistical significance to select discriminative subgraphs. It defines a new measurement counting the number of features that

Fig. 5 The overview of our approach. The red boxes indicate the main contributions of our method.

are not helpful for classification among candidate features. This measurement can be integrated into gSpan as a culling method. It can reduce the number of features, while preserving performance in classification and can prune search space without relying upon a manually-tuned frequency threshold.

A near-optimality of CORK is obtained from a submodular quality function, $q(\cdot)$, using a greedy forward feature selection. The function $q(\cdot)$ considers presence or absence of each subgraph in each class. $q(\cdot)$ for the set containing subgraph, S , is defined as the following:

$$q(\{S\}) = -(A_{S_0} \cdot B_{S_0} + A_{S_1} \cdot B_{S_1}) \quad (2)$$

where A and B are two classes in a dataset. A_{S_0} is the number of images of the class A that do not have the subgraph set $\{S\}$. A_{S_1} is the number of images including the subgraph set $\{S\}$ in the class A . The subscripts S_0 and S_1 are used in the same manner for another class B .

When a subgraph appears or does not appear simultaneously in both classes, it can be considered as a non-discriminative feature between two classes. To consider this observation, A_{S_0} and B_{S_0} are multiplied together; the same reasoning applies for the product of A_{S_1} and B_{S_1} . In this context, a feature becomes more discriminative, as the quality function $q(\cdot)$ becomes higher. Figure 6 shows examples of the quality function for two subgraphs in classes A and B .

While generating subgraphs, we commonly expand a subgraph S into another one, T , by adding a neighboring edge or so. During this incremental process, suppose that we already decided to include S into our feature set. We then need to check the quality of having a newly expanded feature T on top of S . As a result, we need to reevaluate $q(\{T\})$.

Unfortunately, this process can require an

		Dataset					
		Class A			Class B		
Subgraph	S	a_1	a_2	a_3	b_1	b_2	b_3
	T	1	1	1	1	1	0
		1	1	1	0	0	0

Fig. 6 A and B are two different classes in a given dataset. a_1 – a_3 and b_1 – b_3 are images in classes A and B , respectively. Each indicator is 1, if its corresponding subgraph appears in each image, otherwise 0. Referred to Eq. (2), $q(\{S\}) = -(0 \cdot 1 + 3 \cdot 2) = -6$ and $q(\{T\}) = -(0 \cdot 3 + 3 \cdot 0) = 0$. As a result, the subgraph T has a higher discriminative power than S .

excessive amount of running time, since as the number of features increases to N , the number of possible feature combinations can increase exponentially to 2^N .

To accelerate this process, CORK relies on a pruning criterion. Especially, the upper bound of the quality function is derived based on three possible cases, when we consider a supergraph T from its subgraph S . One of such cases is that images from class A do not have the supergraph T , while images in the other class have the supergraph and thus their indicator values are affected. The second case is that the scenario of the first case is applied in the reverse way to classes A and B . The third case is where we do not have any changes. By considering these three different cases, the upper bound of the quality function is derived as the following [26, Theorems 2.2, 2.3]:

$$q(\{T\}) \leq q(\{S\}) + \max \begin{cases} A_{S_1} \cdot (B_{S_1} - B_{S_0}) \\ (A_{S_1} - A_{S_0}) \cdot B_{S_1} \\ 0 \end{cases} \quad (3)$$

While expanding subgraphs, we prune the children of supergraphs T expanded from the subgraph S , when the quality function of T is equal to the one of those upper bounds. This culling criterion is adopted, since it is guaranteed that we cannot find any better supergraphs than T whose quality function is higher than the upper bound shown in the aforementioned inequality. This approach has been proven to identify discriminative subgraphs whose quality function values are bigger than a certain lower bound [26, Theorem 2.1]. Furthermore, unlike gSpan, users do not need to provide manually-tuned parameters for identifying discriminative subgraphs.

4.3 TF-IDF normalization

Once we extract features, we normalize those features. TF-IDF [27] is one of commonly adopted normalization schemes, mainly for document classification. We apply this normalization to our feature, which resembles the bag-of-word model. Inspired by the TF-IDF normalization scheme, we give higher weights to more frequent features in each image and deemphasize features that appear in more images.

In particular, our TF-IDF weighting scheme of a subgraph s occurring in an image i given an image

database D is defined as the following:

$$\begin{cases} TF-IDF(s, i, D) = TF(s, i) \times IDF(s, D) \\ TF(s, i) = \log(1 + f_{s,i}) \\ IDF(s, D) = \log\left(\frac{N}{1 + n_s}\right) \end{cases} \quad (4)$$

where $f_{s,i}$ is the number of the subgraph s occurring in the image i , N is the number of all images in the database D , and n_s is the number of images with the subgraph s . If $f_{s,i}$ is zero, TF term would be undefined. To prevent this case, a small constant, 1, is added. Similarly, to avoid divide-by-zero, we also add the small constant 1 to the denominator of the IDF term.

5 Results

We implemented prior and our methods for discovering family photos in a machine that has Xeon 3.47 GHz with 192 GB main memory. We evaluate the effectiveness of computing and using discriminative feature selection along with TF-IDF normalization. For classification, we use the support vector machine (SVM). The classification is conducted with linear kernel and 5-fold cross validation.

5.1 Datasets

To validate our approach, we use the existing dataset provided by Chen et al. [2]. We additionally test different methods against a new, larger, and diverse dataset, which is rearranged from the public dataset [3], as adopted also in the previous work. Based on the protocol laid out in the prior work, we obtain a “soft” ground truth containing 1613 family photos and 1890 non-family photos for our new, extended dataset. The “soft” ground truth for the new dataset is generated without any prior knowledge such as looking at labels of those images.

The new extended dataset also shares the same images to the Chen et al.’s dataset, since these two datasets are arranged from the public dataset. We also measure the common images in both or either one of two datasets (Table 1). The difference from the previous one is that the new dataset has 1073 more photos and includes wider sets of family types such as siblings, single parent, nuclear family, and extended family, as shown in Fig. 1.

Note that these images from the public dataset

Table 1 The composition of Chen et al.’s ((a)+(b)) and our datasets ((b)+(c)). (b) indicates the number of co-occurring images in both Chen et al.’s and ours. Many images of family and non-family co-occur in ours and Chen et al.’s, although we prepare the extended dataset without looking at their original labels

	(a) Images only in Chen et al.’s	(b) Images both in Chen et al.’s and ours	(c) Images only in ours
Family	66	1111	502
Non-family	136	1131	759

have labels, which are groups, wedding, and family types. Our methods independently predict family types of these images and measure accuracy by comparing their predicted labels with the original ones associated with the public dataset.

We have also considered other datasets related to group photos [4, 8]. Unfortunately, these datasets do not contain labels directly for family and non-family types. As a result, we were unable to use them for our problem.

5.2 Effects of discriminative subgraphs

We test accuracy of different methods including ours and the gSpan method [2]. We have implemented the prior method by following the guideline of the original paper [2]. For gSpan, we generate frequent subgraphs up to 10,000 subgraphs and sort them in the order of document frequency and select them as BoFG. To achieve the best accuracy for the gSpan method, it is required for users to specify the number of subgraphs. In this approach, we need to rely on many trial-and-error procedures, while our method automatically constructs a set of discriminative subgraphs.

We were unclear how the prior method uses the document frequency (DF) term, because there is an ambiguity in which the DF term is evaluated either after or during the running process of gSpan¹. We thus experiment both cases. $gSpan + DF(1)$ and $gSpan + DF(2)$ correspond respectively to the adaption of DF *posterior to* and *during* gSpan in Table 2. In Table 2, our method finds the maximal number of subgraphs without using the minimum frequency.

Our methods w/ and w/o the TF-IDF scheme in the Chen et al.’s dataset identify a small set of discriminative subgraphs (i.e., 76 subgraphs), and achieves 80.61% and 78.65% accuracy respectively.

¹ We have consulted authors of the gSpan technique for faithful re-implementation of the gSpan method.

Table 2 The accuracy of different methods in Chen et al.'s and our datasets

Chen et al.'s dataset									
Dimension	76	100	500	1000	2000	3000	4000	5000	10,000
$gSpan + DF(1)$	50.00%	51.51%	52.98%	54.65%	61.92%	68.33%	69.52%	68.78%	77.76%
$gSpan + DF(2)$	78.61%	77.92%	80.12%	78.16%	77.51%	77.31%	76.49%	77.14%	77.76%
$Ours + TF-IDF$	80.61%	N/A							
Our dataset									
Dimension	90	100	500	1000	2000	3000	4000	5000	10,000
$gSpan + DF(1)$	56.00%	58.37%	62.25%	61.26%	64.51%	67.48%	69.64%	71.84%	75.61%
$gSpan + DF(2)$	74.78%	74.84%	77.43%	76.80%	76.63%	76.83%	76.49%	76.09%	75.61%
$Ours + TF-IDF$	79.34%	N/A							

Table 3 The accuracy of DF vs. $TF-IDF$ in Chen et al.'s and our extended datasets

Chen et al.'s dataset						Our dataset				
Dimension	76	100	1000	5000	10,000	90	100	1000	5000	10,000
$gSpan + DF(2)$	78.61%	77.92%	78.16%	77.14%	77.76%	74.78%	74.84%	76.8%	76.09%	75.61%
$gSpan + DF(2) + \mathbf{TF-IDF}$	77.67%	77.63%	81.31%	81.18%	82.04%	75.40%	75.09%	78.09%	77.55%	77.2%
<i>Ours</i>	78.65%	N/A				77.26%	N/A			
Ours + TF-IDF	80.61%	N/A				79.34%	N/A			

Our method in the extended dataset achieves 77.26%, and shows 79.34% with the TF-IDF scheme. $gSpan + DF(1)$ and $gSpan + DF(2)$ methods show inferior results over our method in most cases. Interestingly, the prior methods show even lower accuracy as they use higher dimensions. This is mainly because frequent subgraphs may not be discriminative.

5.3 Effects of TF-IDF normalization

We measure accuracy of different methods with and without TF-IDF normalization. Since $gSpan + DF(2)$ achieves higher accuracy than $gSpan + DF(1)$, we show the results of $gSpan + DF(2)$ and ours for the test.

In both $gSpan + DF(2)$ and ours, using TF-IDF over DF improves the classification accuracy in most cases. Especially, our method using TF-IDF achieves the highest accuracy, 79.34%, for the extended dataset.

5.4 Comparison of subgraphs

We check the number of subgraphs co-occurring in

the BoFG features generated by both $gSpan$ and our method. This investigation can help us to understand how many dimensions prior methods require in order to obtain discriminative features extracted by our method. Even in hundreds of thousands of dimensions extracted by $gSpan$, some of discriminative subgraphs extracted by our method are not identified (Table 4).

We also measure how well query images used in the test phase are represented by extracted features. For this, we measure how many query images are represented by null vector, indicating that query images are not represented by any features extracted by $gSpan$ or our method (Table 5). As a result, we can conclude that the feature extraction of our method performs better than other tested methods ($gSpan + DF(1)$ or $gSpan + DF(2)$).

6 Conclusions and future work

We have proposed a novel classification system utilizing discriminative subgraph mining for

Table 4 The number of common subgraphs between $gSpan$ and ours in Chen et al.'s and our datasets

Chen et al.'s dataset: 77 subgraphs by our method									
$gSpan$ subgraphs	78	100	500	1000	10,000	...	111,764	...	560,177
Number of common subgraphs	16	17	21	23	36	...	54	...	59
Our dataset: 85 subgraphs by our method									
$gSpan$ subgraphs	85	100	500	1000	10,000	...	84,713	...	326,034
Number of common subgraphs	20	21	25	26	29	...	50	...	59

Table 5 This table shows the number of query images that are represented by the null vector in Chen et al.'s and our extended data

Dimension	76	100	1000	5000	10,000
$gSpan + DF(1)$	283	283	189	82	30
$gSpan + DF(2)$	53	52	39	37	30
Ours	26		N/A		
Dimension	90	100	1000	5000	10,000
$gSpan + DF(1)$	351	351	236	163	43
$gSpan + DF(2)$	62	62	50	48	43
Ours	28		N/A		

achieving high accuracy. We represent group photos as graphs with age, gender, and face position, and then extract discriminative subgraphs and construct BoFG features. For extracting discriminative subgraphs, we proposed to use a recent discriminative subgraph mining method, CORK, that adopts a quality function with near-optimal guarantees. We additionally proposed to use the TF-IDF normalization to better support the characteristic of BoFG features. To validate benefits of our approach, we have tested different methods including ours against two different datasets including our new, extended dataset. Our method achieves higher accuracy in the same dimensionality over the prior methods. Furthermore, our method achieves higher or similar accuracy over the prior work that relies on manual turning and requires a higher dimensionality.

There are many interesting future directions. Since our work is based on the concept of social relationships, we consider subgraphs consisting of at least two nodes. However, only a single node can provide useful social cues. Incorporating single nodes in BoFGs and investigating its effects should be interesting. We would like to also investigate recent deep learning techniques that learns low-level features and classification functions. Due to the lack of sufficient training datasets, we did not consider recent deep learning techniques, but this approach should be worthwhile for achieving higher accuracy.

Acknowledgements

We are thankful to our lab members for valuable feedbacks, and to Ph.D. Yan-Ying Chen for sharing her dataset. This work was supported in part by MSIP/IITP (Nos. R0126-16-1108, R0101-16-0176) and MSIP/NRF (No. 2013-067321).

References

- [1] Krizhevsky, A. Learning multiple layers of features from tiny images. Technical Report. University of Toronto, 2009.
- [2] Chen, Y.-Y.; Hsu, W. H.; Liao, H.-Y. M. Discovering informative social subgraphs and predicting pairwise relationships from group photos. In: Proceedings of the 20th ACM International Conference on Multimedia, 669–678, 2012.
- [3] Gallagher, A. C.; Chen, T. Understanding images of groups of people. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 256–263, 2009.
- [4] Murillo, A. C.; Kwak, I. S.; Bourdev, L.; Kriegman, D.; Belongie, S. Urban tribes: Analyzing group photos from a social perspective. In: Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, 28–35, 2012.
- [5] Wang, G.; Gallagher, A.; Luo, J.; Forsyth, D. Seeing people in social context: Recognizing people and social relationships. In: *Lecture Notes in Computer Science, Vol. 6315*. Daniilidis, K.; Maragos, P.; Paragios, N. Eds. Springer Berlin Heidelberg, 169–182, 2010.
- [6] Chiu, Y.-I.; Li, C.; Huang, C.-R.; Chung, P.-C.; Chen, T. Efficient graph based spatial face context representation and matching. In: Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing, 2001–2005, 2013.
- [7] Gallagher, A. C.; Chen, T. Finding rows of people in group images. In: Proceedings of IEEE International Conference on Multimedia and Expo, 602–605, 2009.
- [8] Choi, W.; Chao, Y.-W.; Pantofaru, C.; Savarese, S. Discovering groups of people in images. In: *Lecture Notes in Computer Science, Vol. 8692*. Fleet, D.; Pajdla, T.; Schiele, B.; Tuytelaars, T. Eds. Springer International Publishing, 417–433, 2014.
- [9] Shu, H.; Gallagher, A.; Chen, H.; Chen, T. Face-graph matching for classifying groups of people. In: Proceedings of IEEE International Conference on Image Processing, 2425–2429, 2013.
- [10] Chiu, Y.-I.; Hsu, R.-Y.; Huang, C.-R. Spatial face context with gender information for group photo similarity assessment. In: Proceedings of the 22nd International Conference on Pattern Recognition, 2673–2678, 2014.
- [11] Shimizu, K.; Nitta, N.; Nakai, Y.; Babaguchi, N. Classification based group photo retrieval with bag of people features. In: Proceedings of the 2nd ACM International Conference on Multimedia Retrieval, Article No. 6, 2012.
- [12] Zhang, T.; Chao, H.; Willis, C.; Tretter, D. Consumer image retrieval by estimating relation tree from family photo collections. In: Proceedings of the ACM International Conference on Image and Video Retrieval, 143–150, 2010.
- [13] Singla, P.; Kautz, H.; Luo, J.; Gallagher, A. Discovery of social relationships in consumer photo

- collections using Markov logic. In: Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, 1–7, 2008.
- [14] Song, Z.; Wang, M.; Hua, X.-s.; Yan, S. Predicting occupation via human clothing and contexts. In: Proceedings of International Conference on Computer Vision, 1084–1091, 2011.
- [15] Lowe, D. G. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* Vol. 60, No. 2, 91–110, 2004.
- [16] Oliva, A.; Torralba, A. Modeling the shape of the scene: A holistic representation of the spatial envelope. *International Journal of Computer Vision* Vol. 42, No. 3, 145–175, 2001.
- [17] Jiang, C.; Coenen, F.; Zito, M. A survey of frequent subgraph mining algorithms. *The Knowledge Engineering Review* Vol. 28, No. 1, 75–105, 2013.
- [18] Cook, S. A. The complexity of theorem-proving procedures. In: Proceedings of the 3rd Annual ACM Symposium on Theory of Computing, 151–158, 1971.
- [19] Wang, M.; Lai, Y.-K.; Liang, Y.; Martin, R. R.; Hu, S.-M. BiggerPicture: Data-driven image extrapolation using graph matching. *ACM Transactions on Graphics* Vol. 33, No. 6, Article No. 173, 2014.
- [20] Chen, X.; Zhou, B.; Guo, Y.; Xu, F.; Zhao, Q. Structure guided texture inpainting through multi-scale patches and global optimization for image completion. *Science China Information Sciences* Vol. 57, No. 1, 1–16, 2014.
- [21] Li, H.; Wu, W.; Wu, E. Robust interactive image segmentation via graph-based manifold ranking. *Computational Visual Media* Vol. 1, No. 3, 183–195, 2015.
- [22] Hu, S.-M.; Zhang, F.-L.; Wang, M.; Martin, R. R.; Wang, J. PatchNet: A patch-based image representation for interactive library-driven image editing. *ACM Transactions on Graphics* Vol. 32, No. 6, Article No. 196, 2013.
- [23] Taigman, Y.; Yang, M.; Ranzato, M.; Wolf, L. DeepFace: Closing the gap to human-level performance in face verification. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, 1701–1708, 2014.
- [24] Zhao, W.; Chellappa, R.; Phillips, P. J.; Rosenfeld, A. Face recognition: A literature survey. *ACM Computing Surveys* Vol. 35, No. 4, 399–458, 2003.
- [25] Yan, X.; Han, J. gSpan: Graph-based substructure pattern mining. In: Proceedings of the 2002 IEEE International Conference on Data Mining, 721–724, 2002.
- [26] Thoma, M.; Cheng, H.; Gretton, A.; Han, J.; Kriegel, H.-P.; Smola, A.; Song, L.; Yu, P. S.; Yan, X.; Borgwardt, K. M. Discriminative frequent subgraph mining with optimality guarantees. *Statistical Analysis and Data Mining: The ASA Data Science Journal* Vol. 3, No. 5, 302–318, 2010.
- [27] Jones, K. S. A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation* Vol. 28, No. 1, 11–21, 1972.

Changmin Choi is in a start-up company. He received his M.S. degree from School of Computing at Korea Advanced Institute of Science and Technology (KAIST), and B.A. degree from Business School at Hanyang University. His research interest is understanding group photos in social

media.

YoonSeok Lee is a M.S. student in School of Computing at Korea Advanced Institute of Science and Technology (KAIST) and he received his B.S. degree in computer science from KAIST in 2014. His research interest lies in image classification, image representation, and hashing techniques.

Sung-Eui Yoon is currently an associate professor at Korea Advanced Institute of Science and Technology (KAIST). He received his B.S. and M.S. degrees in computer science from Seoul National University in 1999 and 2001, respectively. His main research interest is designing scalable graphics,

image search, and geometric algorithms. He gave numerous tutorials on proximity queries and large-scale rendering at various conferences including ACM SIGGRAPH and IEEE Visualization. Some of his work received a distinguished paper award at Pacific Graphics, invitations to IEEE TVCG, an ACM student research competition award, and other domestic research-related awards. He is a senior member of IEEE, and a member of ACM and KIISE.

Open Access The articles published in this journal are distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Other papers from this open access journal are available free of charge from <http://www.springer.com/journal/41095>. To submit a manuscript, please go to <https://www.editorialmanager.com/cvmj>.