Aiding Face Recognition with Social Context Association Rule based Re-Ranking

Samarth Bharadwaj, Mayank Vatsa, and Richa Singh
IIIT-Delhi, India
{samarthb, mayank, rsingh}@iiitd.ac.in

Abstract

Humans are very efficient at recognizing familiar face images even in challenging conditions. One reason for such capabilities is the ability to understand social context between individuals. Sometimes the identity of the person in a photo can be inferred based on the identity of other persons in the same photo, when some social context between them is known. This research presents an algorithm to utilize co-occurrence of individuals as the social context to improve face recognition. Association rule mining is utilized to infer multi-level social context among subjects from a large repository of social transactions. The results are demonstrated on the G-album and on the SN-collection pertaining to 4675 identities prepared by the authors from a social networking website. The results show that association rules extracted from social context can be used to augment face recognition and improve the identification performance.

1. Introduction

Face recognition capabilities of humans have inspired several researchers to understand the *science* behind it and use it in developing automated algorithms. Recently, it is also argued that encoding *social context* among individuals can be leveraged for improved automatic face recognition [26]. As shown in Figure 1, often times a person's identity can be *inferred* based on the identity of other persons in the same photo, when some social context between them is known. A subject's face in consumer photos generally co-occur along with their socially relevant people. With the advent of social networking services, the social context between individuals is readily available. Face recognition performance in such photos can be improved by considering this contextual information and can have interesting applications in image sharing, forensics, and intelligent surveillance systems.

The term *context* has been used in object recognition [27], person detection and also in face recognition research [17, 24] to imply acceptable co-occurrence of various parts or attributes of an object or face. Researchers compute

Figure 1. In this illustration, context derived from two images help confirm identity of a face in third. *A professor meets a couple of PhD students, A and T, after an invited talk at a university. Later at a conference, he is unable to recognize A but infers his identity when he sees A and T seated together.*

context from scene recognition and/or spatiotemporal information obtained from image headers, camera devices, user labels for events, locations, people, and other secondary sources. In this research, *social context* refers to the possibility of co-occurrence of two or more individuals in a single photo. Since consumer photography is generally a social exercise, involving people of social importance, an improved understanding of *social context* can augment the quality of face recognition and tagging.

1.1. Related Work

We briefly review literature of various problem domains that use social context to aid face recognition performance. As illustrated in Figure 2, augmenting the performance of face recognition using various cues have been studied in literature with two application scenarios:
i) *Organizing photo album collections*: Personal photo collections stored in computers and smart phones are generally clustered based on events, people, locations. Several research directions have been undertaken to improve automatic photo organizer by attaching meaningful labels pertaining to identities, relationships, and other demographics including the use of meta-information from capture devices (cellID, GPS, time) [8, 21]. The popularity of digital photos has also lead towards commercial photo management tools

(a) (b) (c)

Figure 2. Three problem domains: (a) traditional face recognition, (b) photo album organization (based on events, locations and people) with imperfect (manual) annotations and (c) reliable face tagging approaches of large number of uploaded photos, which is the focus of this work.

such as Google Picasa, Microsoft EasyAlbum [6] and Apple iPhoto.

Kapoor *et al*. [16] used active learning to minimize manual labor by inferring probabilistic models to simultaneously tag people, events, and locations by deriving cross-domain relationships in semi-supervised settings. Other researchers have used heuristically generated priors based on rules such as *height of husband is greater than height of wife* [28]. O'Hare *et al*. [21] combined several context cues derived from text, event detection, image color descriptors and body part analysis to improve person identification in photo collections. Gallagher and Chen [12] improved ambiguous person labels in image collection by learning group priors to classify unlabeled persons. In another approach [11], they evaluated relative position of people in group photos to improve gender and age classification. Further, clothing [10] is also used to cluster images from the same event. Face clustering approaches incorporate contextual constraints such as *same-day* (a person on a given day is wearing the same clothing throughout), *Person-Exclusive* constraint (a person can occur in a photo only once), and co-occurrence [30] to improve recall. Manyam *et al*. [20] questioned the independence assumption between features extracted from two faces in the same image. Exploiting the consistent environmental settings, relative features are extracted to augment joint face recognition with a conditional probabilistic model. Chen *et al*. [5] infer pair-wise relationships and social-subgroups within images (such as siblings, couple etc.) based on term frequency of low level visual features. Satish *et al*. [23] explore context discovery from multiple information sources. Recently, Barr *et al*. [2] use active learning to generate social constraints to improve face clustering and Hochreiter *et al*. [14] model co-occurrence of individuals in photo albums.

ii) *Reliable face tagging*: It has extensive utility, particularly for cloud based storage and online social networking services. A large number of consumer images are stored for easy access, sharing, and reliable storage (with over 300 million photos per day uploaded on Facebook alone [15]). The shared and tagged photos provide extensive social context to predict possible labels of an input photo without manual intervention. Becker *et al*. [3] perform preliminary evaluation of several established appearance based approaches on a set of images mined from Facebook. In another study, Stone *et al*. [25, 26] perform a large scale face recognition study on photos from Facebook. The match scores obtained between two subjects are augmented with a social context score, that measures whether both subjects are *friends* (acquaintance of one another on the social network site) or not. Hence, the social cue corrects the face match score and is suitably weighted within a conditional random field framework. The approach is also implemented on mobile phones [7], showcasing the applicability of social context augmentation in face recognition in realtime applications. Lin *et al*. [18] present a framework of tagging face images in consumer photos. Recently, Wu *et al*. [29] use belief propagation for identity discovery in a synthetically generated social network. Sapkota *et al*. [22] use an ensemble of context and facial feature classifiers to improve face recognition performance.

1.2. Research Contributions

As discussed, several techniques have been proposed in literature to augment face recognition with a social context.

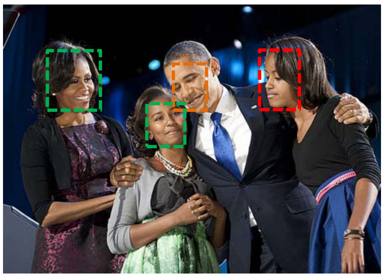

Figure 3. In case of challenging imaging condition, prior knowledge of association of subjects may induce additional information to improve face recognition. The proposed approach generates context association rules, $\{Michelle, Sasha\} \rightarrow Barack$ and $\{Michelle, Barack, Sasha\} \rightarrow Malia$, based on prior co-occurrence to aid face recognition.

Social context such as clothing color, relative location of faces, and age/gender estimation have shown promising results on consumer photo datasets of relatively small size. In such cases, the number of subjects are limited to 10-15 (for example, family members). In order to robustly use context information in large, more uncertain consumer photo collections, such as social networking sites or surveillance systems, the applicability of context must be more closely re-examined.

This research aims to broaden the scope of face recognition in consumer photos using social context. Rather than binary cues that have been explored in literature, such as $\{friend, no{-}friend\}$, the proposed approach infers association between groups of individuals from multi-level social cues such as co-occurrence of people in consumer photos, to improve face identification. These context cues are used to re-rank face recognition results to improve the overall performance. Secondly, to evaluate the performance of the proposed algorithm, a large dataset is mined from a leading social networking site consisting of 160,264 images from 4675 connected users. An anonymized subset of the social graph, together with rank orders of identities obtained from a powerful face recognition system on the uploaded photos are made available to the research community.

2. Social Context Aided Face Recognition

The key concept of the proposed algorithm is illustrated in Figure 3. The algorithm first learns the social context from a set of training images. The contextual information is used to improve the face recognition ranking of a probe image obtained from any face recognition system. To replicate a practical scenario where a user uploads several consumer photos to a photo sharing service, a non-overlapping face identification scenario is considered in this research. The social context obtained is used to augment face recognition in unseen faces. The details of the proposed approach are as follows.

2.1. Building Context from Tagged Faces

In this research, co-occurrence of people in a consumer photo is used as transactions to improve face recognition performance. By viewing a photo consisting of more than one individual as a single *social transaction*, it is possible to infer contextual association rules (AR) that capture social context. For instance, a rule $\{Michelle, Barack, Sasha\} \rightarrow Malia$, can assist in recognition of off-angle faces in challenging imaging conditions such as in Figure 3. Association rules of the form $X \rightarrow Y$ are mined from *tagged* faces in group photos. Since there is no restriction on the cardinality of the antecedent of rule (X), the social cue inherently provides richer information regarding the co-occurrence of subjects. While tagging may not be reliable for accurate face labeling, useful contextual information may be derived from each image, based on which users co-occur in tags.

During the training phase, a large number of *social transactions* are utilized to create social context association rules. Higher frequency of occurrence of the same set of individuals in a large collection of social transactions is inferred as *a greater chance of their co-occurrence in the probe images*. Association rule mining is used to derive inferences from the training samples. Several such association rules encode the hierarchical (level-wise) co-occurrence between different clusters of individuals. In this work, we associate photos that contain two or more *tagged* individuals as a single social transaction. Further, we hypothesize that association rules thus derived can be used to augment face recognition performance in challenging settings.

In the testing phase illustrated in Figure 4, face detection and recognition are first applied to each probe image and rank ordering of gallery for each of the detected faces is obtained. Next, based on the confidence of detection and recognition, social context rules involving the most confident match are used to *influence* the rank ordering of the remaining faces. Finally, the set of rank orders obtained via face recognition and different combinations of social context are fused to obtain a single rank order. The proposed approach is presented in Algorithm 1. A two stage augmentation is applied to each probe face, described as follows:

- **Unique Identity Pruning:** It can be safely assumed that the same person does not occur in a single photo more than once. Hence, if a label is assigned confidently to a face, the probability of the same label being assigned to another face in the same photo is suitably reduced.
- **Social Co-occurrence:** Based on the tags from previous photos in the collections, multi-level social context is inferred and used to improve face recognition by adjusting rank order. The mining of association rule and fusion scheme are discussed in detail next.

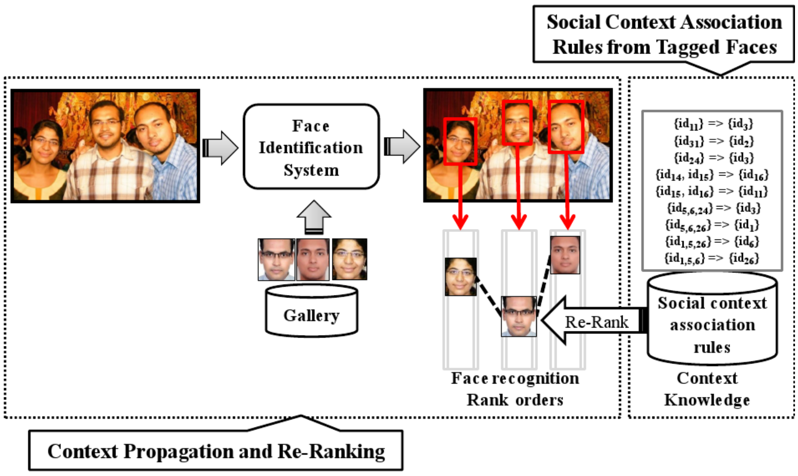

Figure 4. An illustration of the proposed multi-level association rule mining to derive social context to improve face identification.

2.2. Social Context Association Rules from Tagged Faces

Association rule mining is a data mining technique widely used in literature to extract rules from transactional data for applications such as market-basket analysis. These rules provide insights into the behavior of customers by inferring relationships between frequent items that are in consideration. AR mining is also used in web user behavior analysis, intrusion detection, DNA sequencing, various web recommendation services and sub-group discovery. Liu *et al*. [19] use the Apriori algorithm [1] to post-filter semantic concepts that are detected in videos using association rules between known semantic concepts. An association rule of the concept *sky→outdoor* has higher confidence than with *indoor*, hence post-processing improved performance of concept detection in videos. Given a set of social transactions, we define the following:

- **Social Context Association Rule**: An association rule of the form X→Y represents a social context between two *sets* of subject labels X and Y. Here, X∩Y = ∅ and both X and Y are subsets of U, the set of all face labels. For multi-level context mining, only those association rules that have one item as consequent (righthand side) are considered, i.e., $\{I_1, I_2, \ldots I_{k-1}\} \rightarrow I_k$ for a set of tagged users, $I_1, I_2, \ldots I_k$. These type of rules are referred to as Class Association Rules (CAR).

- **Support** of a set X denoted by **Supp**(X), is the probability of occurrence of the members of X in the given set of transactions. The support of an association rule of the form X→Y represents the probability that members of X and Y have co-occurred in the transactions.

- **Confidence** of an association rule represents the probability that members of X and Y have co-occurred given the occurrence of X. Confidence may also be interpreted as the conditional probability $P(Y|X)$.

We consider the extensively used association rule mining algorithm, Apriori [1], that uses breadth-first (level-wise) search to determine rules based on the downward closure property of support. The iterative approach effectively finds frequent sets of k subjects that co-occur in the social transactions. Further, k subject set is used to generate $(k+1)$ frequent subject set, focusing only on subjects that frequently occur in the transactions. Further, a pre-defined minimum threshold, min_suppt, is maintained to make the breadth-first search computationally tractable. The association rules are stored in lookup table (ApR) with the rule confidence (C) as *data* and rule items $(\{X \cup Y\})$ as the *key*. For large datasets, rules with higher support can be cached in memory, since they are more expensive to re-compute. In comparison to existing graphical modeling approaches in literature, association rules inherently capture a higher cardinality of associations. Figure 5 shows a set of 100 rules connected based on confidence. As illustrated, the rules capture dense information pertaining to n-tuple co-occurrences.

2.3. Context Propagation and Re-Ranking

Face recognition based rank order is first evaluated for confidence. The confident face match is used as evidence to infer contextual information. To compute the confidence of a rank order, the likelihood ratio $\frac{P(X=l_1)}{P(X=l_2)}$ of scores between top match and the next closet match is considered. A high likelihood ratio indicates a confident match that forms the basis for deriving social context of other faces in the same photo. The proposed approach builds social context in a photo starting from the *most confident* face match. Given a confident match x_{conf}, unique identity pruning is first performed on the remaining faces, as shown in Eq. 1.

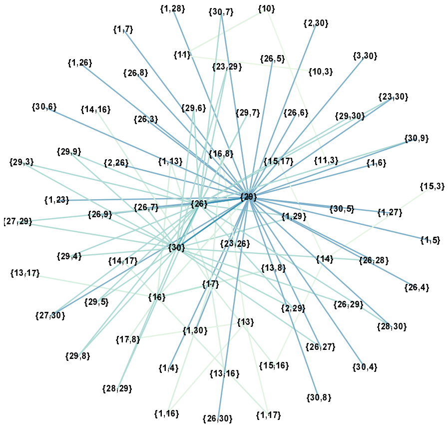

Figure 5. A visualization (Fruchterman Reingold layout) of 100 association rules [13] with the highest *lift* obtained from the G-album dataset [10]. Lift is defined as $\frac{\textbf{Supp}(X \cup Y)}{\textbf{Supp}(X)\textbf{Supp}(Y)}$, greater lift is indicative of more informative association rules. The vertices are shaded proportional to the in-degree of each rule, i.e., the number of rules the consequent participates in.

$$\hat{P}(x) = (1 - P_{face}(x_{conf})) \times P_{face}(x) \qquad (1)$$

Next, the rank ordering obtained from a face recognition system is re-ranked based on the social association rules obtained in the training phase. The updated score for a particular identity ($R(x)$) after re-ranking is obtained by a weighted aggregation of the normalized similarity scores obtained from face recognition and contextual information (confidence of association rules).

$$R(x) = \alpha P_{face}(x) + (1-\alpha)\left[\frac{1}{2^n}\sum_{T \in \mathcal{P}(X)} P_{context}(x_n|T)\right] \qquad (2)$$

where $P_{face}(x)$ is obtained from a face recognition system. $P_{context}(Y|X)$ is the confidence of identity Y, given the occurrence of $X = \{F_1, F_2 \ldots F_{n-1}\}$ obtained from the corresponding association rule $X \rightarrow Y$. Further, the confidence values of the power set of X, denoted by $\mathcal{P}(X)$, are combined. The iterative accumulation of evidence of context ensures that all subsets of association are also considered, i.e., both $\{Michelle, Sasha\} \rightarrow Barack$ and $\{Michelle, Barack, Sasha\} \rightarrow Malia$ rules are evaluated sequentially. The confidence value α is a scalar quantity that favors the relevance score provided by face recognition.

3. Database and Protocol

The proposed approach is evaluated on a publicly available database (G-album) [10] along with a large database, SN-collection, collected by the authors from a social networking site. These face databases are selected instead of standard large test databases in order to demonstrate the advantage of derived social context in augmenting face recognition performance on challenging consumer photos. The details of the databases are summarized in Table 1.

3.1. The G-album

It consists of 589 images pertaining to 32 subjects from various family events. The photographs are taken in unconstrained settings with illumination, expression and pose variations. Most of the photographs consist of more than one subject with high co-occurrence of several subjects. The experimental protocol is designed to replicate a typical scenario in which a collection of existing photographs are available with tagging. For the G-album database, 50% images from the entire database are set aside as probe (depicting uploaded images) and the remaining are used as training samples. This forms a seen experiment, with different images of the same subjects in training and testing.

3.2. SN-collection

A major bottleneck in the progression of this research area has been the lack of a large database for comparison and evaluation. Photos from social networking services and corresponding social connections of the users are difficult to share without violating privacy concerns. Hence, for the benefit of the research community, we prepared the SN-collection database. The dataset is prepared using $160,264$ images that are obtained from 65 active users (undergraduate students) and their self-declared, bi-directional *friends*, a total of $4,675$ users. The experimental results are demonstrated on photos that are all tagged. A tag is a manually

Algorithm 1 Proposed context association rule based re-ranking approach

input: Given social transactions for training T, gallery set G recognize faces in probe P
face detection: detect faces F in P
training phase:
create look-up table ApR of association rules from T

testing phase:
compute C_i, the confidence of match for F_i in P
sort: F by C
apply unique identity pruning on rankorders in F(as in Eq. 1)
compute context scores S = Context(F, ApR)
update scores by re-ranking F (as in Eq. 2)
output: social context re-ranked faces for probe P

procedure: Context(F, ApR)
compute $Rankscores$ of F_n with social context from all subsets of $F_1, F_2 \ldots F_{n-1}$
return: $Rankscores$
end procedure:

annotated identity label assigned to each face in a photo. One random face tag that overlaps with face detection rectangle of FaceVacs is assumed to be correctly labeled and used in gallery. In our experiments, 1455 images form the gallery set (containing one face instance per user) and 2893 images form the probe (note that multiple faces co-occur in a single image). From the remaining images, 10433 social transactions are obtained and used to extract social context association rules, at a $min\_suppt = 0.001$. While tags may not be reliable to supplement face recognition or detection, useful contextual information may be derived from each image, based on which users co-occur in tags. The images obtained are extremely challenging and include illumination, expression, and pose variations along with low image resolution images. Another challenge that makes the with social networking data limiting [25] is that there are several subjects that are not under consideration making the problem an open-set experiment.

An anonymized social graph of this collection will be made available to the research community[1]. As part of the distribution, we provide tagged face IDs of individuals (in anonymous form) within the social group in consideration. Further, we provide match scores obtained from a commercial face recognition system, FaceVacs. Due to privacy and legal constraints, face images or any other identifiable information are not provided. It is our assertion that this dataset will allow researchers to examine the social context in the photos and improve automatic tagging that is assisted with face recognition technology.

	The G-album	SN-collection
Probe samples	524	2893
No. of subjects	32	4675
No. of detected faces	971	8316
No. of transactions	65	10433
No. of association rules	223312	72455

Table 1. Experimental protocol of both the databases.

4. Experimental Results and Analysis

The proposed social context based face recognition approach is presented in identification ($1 : N$) settings, as it is a more realistic application scenario. Any vanilla face recognition that provides a rank ordered gallery set can be augmented with social context extracted using the proposed approach. Here, FaceVacs face recognition system is used for face detection and matching. The system provides similarity scores between pairs of face images that are used to generate rank-ordering of the gallery set.

- The context association rules are mined from the social transactions from the G-album with min_suppt of 0.01 and 0.001 for SN-collection owing to the size of each dataset. A low support is indicative of low chance of occurrence of a set of individuals in the entire dataset. Further, for both datasets, only rules with confidence greater than 0.5 are chosen. A high confidence indicates high chance of their co-occurrence. On Intel Quad Core (2.5 GHz) processor and 4GB RAM, the training time on G-album is 23.4 seconds and SN-collection is 160.3 seconds.

- With sufficient number of transactions, rules with high confidence can be generated. As shown in Figure 6, the mean confidence of association rules obtained at higher cardinality are more confident. It implies that association rules obtained due to frequent co-occurrence of a large number of individuals are more confident.

- It has been reported in literature that annotation of faces in social networking sites are not reliable enough to be used as face detection [25]. In the proposed approach the tagged faces are used only to derive social context rather than to obtain facial features for matching. We affirm that the social context of co-occurrence of individuals is more likely to be maintained even if the tagging of facial regions in not positionally accurate.

- Figure 7 shows the identification accuracy (using cumulative match characteristics (CMC) curve) of only the face recognition system and the improvement achieved by multi-level social context. The improvement in performance is attributed to the re-ranking based on association rules obtained from the train-

[1]The database will be available at `https://research.iiitd.edu.in/groups/iab/`.

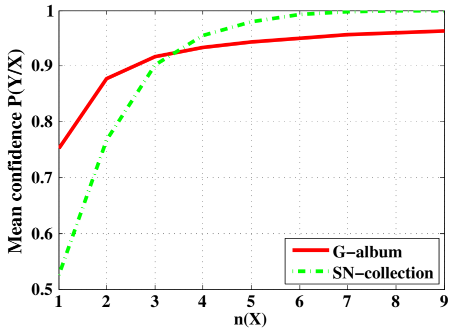

Figure 6. Mean confidence vs. cardinality of X of a social association rule ($X \rightarrow Y$) shows more confidence as $n(X)$ increases.

	The G-album (Rank-1)	SN-collection (Rank-25)
Face Only	79.39%	50.65%
Stone *et al.* [25]	79.41%	50.65%
Proposed	81.33%	55.72%

Table 2. Identification accuracy (%) on two databases.

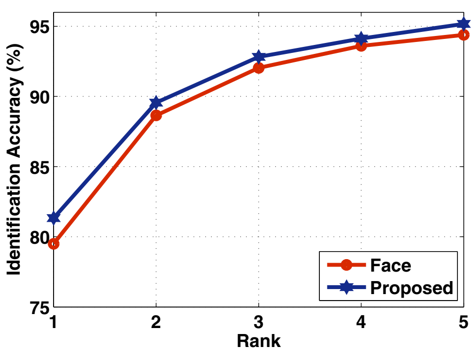

(a)

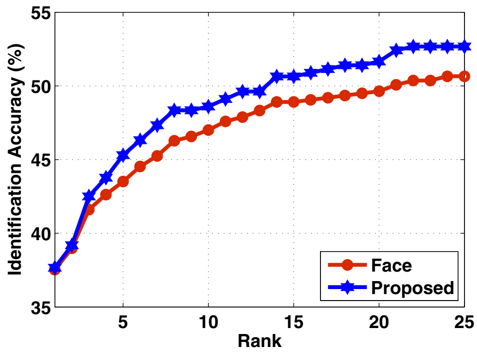

(b)

Figure 7. Identification accuracy obtained on a) G-Album database and b) SN-collection database.

Figure 8. Illustrating three cases from the G-album dataset where initially a face is recognized incorrectly at rank-1. With the application of the proposed context association rule based re-ranking, using the correctly recognized faces and social context, rank-1 performance is achieved for these samples.

ing data. Moreover, it is observed that, on both the datasets, retrieving context information during test/probe has marginal effect on overall computational time. A Kendall tau rank correlation test (social context vs face recognition) shows low correlation: 0.14 with G-album and 0.18 with SN-collection. This indicates that social-context provides additional evidence of identity that improves recognition performance.

- As shown in Table 2, the performance of the proposed approach improves when compared to face recognition alone. Further, the proposed approach is compared with pair-wise conditional random field [25], where pair-wise constraints are obtained by combining co-occurrence and friendship bits. The proposed approach is at an advantage as higher order co-occurrences are utilized which lead to improved accuracy.

- In order to investigate the performance of the re-ranking approach, it is compared with Borda count rank aggregation [9]. This approach performs weighted fusion of rank orders obtained from face recognition and context rather than relevance probabilities. However, the performance of rank fusion approach is found to be lower than the proposed approach. On the G-album, it provides the rank-1 accuracy of 80.33% and on the SN-collection, rank-25 of 50.65% is achieved.

- The performance of the proposed context propagation approach is dependent on the correctness of the most confident face recognition aided tag. In album organization applications, it is often assumed that some faces are manually tagged by users. To explore the performance of context in face recognition for scenario illustrated in Figure 1, rank accuracy with single correct label is computed. The performance of the proposed approach with the assumption of a single correct label improves rank-1 accuracy to 92.5% on G-album and 76.71% on SN-collection. This also indicates that the proposed approach benefits from precise face matching.

- We re-emphasize that the experimental setup is motivated from social networking and cloud photo storage services (illustrated in Figure 2c). The social context is derived from a partial subset of the available datasets and used to predict face labels in *unseen* photos. This is a deviation from existing work (discussed earlier) that focuses on album organization application where context from the entire dataset is used from face clustering.

- Figure 8 illustrates sample instances where one face (right most) is incorrectly matched at rank-1 using face recognition alone. However, the proposed context re-ranking approach uses social context derived from the other images to achieve rank-1 accuracy.

5. Conclusion and Future Work

This research presents an algorithm to augment automatic face recognition with multi-level social context. The proposed approach utilizes association rule mining techniques to extract social cues from co-occurrence of individuals in consumer photos. It is evaluated on G-album, a small publicly available database, and SN-collection, a large database collected by the authors from a social networking site. The results show that multi-level social context helps in improving face identification performance with marginal computational overhead. Social context can be improved with online rule generation and combined with other cues [4] to further improve performance.

References

[1] R. Agrawal and R. Srikant. Fast algorithms for mining association rules. In *VLDB*, volume 1215, pages 487–499, 1994.

[2] J. R. Barr, L. A. Cament, K. W. Bowyer, and P. J. Flynn. Active clustering with ensembles for social structure extraction. In *WACV*, pages 969–976, March 2014.

[3] B. Becker and E. Ortiz. Evaluation of face recognition techniques for application to facebook. In *FG*, pages 1–6, 2008.

[4] S. Bharadwaj, M. Vatsa, and R. Singh. Biometric quality: a review of fingerprint, iris, and face. *JIVP*, 34(1):1–28, 2014.

[5] Y. Y. Chen, W. H. Hsu, and H.-Y. M. Liao. Discovering informative social subgraphs and predicting pairwise relationships from group photos. In *MM*, pages 669–678, 2012.

[6] J. Cui, F. Wen, R. Xiao, Y. Tian, and X. Tang. Easyalbum: an interactive photo annotation system based on face clustering and re-ranking. In *CHI*, pages 367–376, 2007.

[7] M. Dantone, L. Bossard, T. Quack, and L. van Gool. Augmented faces. In *ICCV Workshops*, pages 24–31, 2011.

[8] M. Davis, M. Smith, J. Canny, N. Good, S. King, and R. Janakiraman. Towards context-aware face recognition. In *MM*, pages 483–486, 2005.

[9] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *WWW*, pages 613–622, 2001.

[10] A. Gallagher and T. Chen. Clothing cosegmentation for recognizing people. In *CVPR*, pages 1–8, 2008.

[11] A. Gallagher and T. Chen. Understanding images of groups of people. In *CVPR*, pages 256–263, 2009.

[12] A. C. Gallagher and T. Chen. Using group prior to identify people in consumer images. In *CVPR*, pages 1–8, 2007.

[13] M. Hahsler, S. Chelluboina, K. Hornik, and C. Buchta. The arules r-package ecosystem: Analyzing interesting patterns from large transaction data sets. *JMLR*, 12:2021–2025, 2011.

[14] J. Hochreiter, Z. Han, S. Z. Masood, S. Fonte, and M. Tappen. Exploring album structure for face recognition in online social networks. *IVC*, DOI: 10.1016/j.imavis.2014.01.002, 2014.

[15] http://www.facebook.com/press/info.php?statistics. Facebook statistics[online].

[16] A. Kapoor, G. Hua, A. Akbarzadeh, and S. Baker. Which faces to tag: Adding prior constraints into active learning. In *ICCV*, pages 1058–1065, 2009.

[17] N. Kumar, A. Berg, P. Belhumeur, and S. Nayar. Attribute and simile classifiers for face verification. In *ICCV*, pages 365–372, 2009.

[18] D. Lin, A. Kapoor, G. Hua, and S. Baker. Joint people, event, and location recognition in personal photo collections using cross-domain context. In *ECCV*, pages 243–256, 2010.

[19] K. Liu, M. Weng, C. Tseng, Y. Chuang, and M. Chen. Association and temporal rule mining for post-filtering of semantic concept detection in video. *TMM*, 10(2):240–251, 2008.

[20] O. Manyam, N. Kumar, P. Belhumeur, and D. Kriegman. Two faces are better than one: face recognition in group photographs. In *IJCB*, pages 1–8, 2011.

[21] N. O'Hare and A. Smeaton. Context-aware person identification in personal photo collections. *TMM*, 11(2):220–228, 2009.

[22] A. Sapkota, R. Gopalan, E. Zavesky, and T. Boult. Propagation of facial identities in a social network. In *BTAS*, pages 1–7, 2013.

[23] A. Satish, R. Jain, and A. Gupta. Context networks for annotating personal media [TR]. Technical Report ESL.UCI.EDU-TR, UCI, 2013.

[24] W. Scheirer, N. Kumar, K. Ricanek, P. Belhumeur, and T. Boult. Fusing with context: A bayesian approach to combining descriptive attributes. In *IJCB*, pages 1–8, 2011.

[25] Z. Stone, T. Zickler, and T. Darrell. Autotagging facebook: social network context improves photo annotation. In *CVPR Workshops*, pages 1–8, 2008.

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

[27] A. Torralba, K. Murphy, W. Freeman, and M. Rubin. Context-based vision system for place and object recognition. In *ICCV*, pages 273–280, 2003.

[28] G. Wang, A. Gallagher, J. Luo, and D. Forsyth. Seeing people in social context: Recognizing people and social relationships. In *ECCV*, volume 6315, pages 169–182, 2010.

[29] T. Wu, J. Phillips, and R. Chellappa. Propagation of facial identities in a social network. In *BTAS*, pages 1–7, 2013.

[30] L. Zhang, D. V. Kalashnikov, and S. Mehrotra. A unified framework for context assisted face clustering. In *ICMR*, pages 9–16, 2013.