

On Low-Resolution Face Recognition in the Wild: Comparisons and New Techniques

Pei Li¹ Loreto Prieto² Domingo Mery² Patrick J. Flynn¹

¹Department of Computer Science and Engineering

University of Notre Dame

²Department of Computer Science

Pontificia Universidad Católica de Chile

Abstract—Although face recognition systems have achieved impressive performance in recent years, the low-resolution face recognition task remains challenging, especially when the low-resolution faces are captured under non-ideal conditions, as is common in surveillance-based applications. Faces captured in such conditions are often contaminated by blur, non-uniform lighting, and non-frontal face pose. In this paper, we analyze face recognition techniques using data captured under low-quality conditions in the wild. We provide a comprehensive analysis of experimental results for two of the most important applications in real surveillance applications, and demonstrate practical approaches to handle both cases that show promising performance. The following three contributions are made: (i) we conduct experiments to evaluate super-resolution methods for low-resolution face recognition; (ii) we study face re-identification on various public face datasets including real surveillance and low-resolution subsets of large-scale datasets, present a baseline result for several deep learning based approaches, and improve them by introducing a Generative Adversarial Network (GAN) pre-training approach and fully convolutional architecture; and (iii) we explore low-resolution face identification by employing a state-of-the-art supervised discriminative learning approach. Evaluations are conducted on challenging portions of the SCface and UCCSface datasets.

I. INTRODUCTION

In recent years, we have witnessed tremendous improvements in face recognition system performance, especially when employing different specially designed deep learning architectures. Some state-of-the-art methods based on deep learning models including [66], [60], [51], [67], [77] and [42] have achieved an accuracy over 99% on public face datasets such as LFW [21]. Although these algorithms can deal effectively with faces with significant pose variations, these faces generally need to be large in area. Also, pre-processing techniques such as face frontalization and face alignment are needed. These procedures, that are designed for HR face image data, cannot be applied directly to low-quality face images. The use of surveillance systems in a wide range of public places is increasing, creating a challenging use case for face recognition in an environment where detected faces will be in low-resolution (LR). We refer to this as the *Low-Resolution Face Recognition* (LRFR) problem. Practical face recognition systems for images captured in surveillance scenarios can address face identification tasks (using a watch list) and face

re-identification tasks (where a subject is matched to a previous appearance in a surveillance system).

This work is based in part on prior work published in [37], [36] but presents many new contributions. In the original work [37], we focus on prototyping lightweight deep neural networks for LR face re-identification and showed that LR face images can be used for person re-identification. We employed data from a real surveillance system and conducted an extensive person re-identification study. We address the LR face recognition task in a larger scope which includes a qualitative exploration of the recognition performance gap between controlled facial images and video quality face images using super-resolution techniques, face identification, and face re-identification. We also focus on boosting the performance based on [37] and extend our evaluations using more datasets including SCface, UCCSface and the MegaFace challenge 2 LR subset.

In Section II, we provide a literature review and summary descriptions of several newly released real surveillance datasets. In Section III, we present new methods and experiments as follows: Section III-B presents several baseline face recognition results obtained by evaluating state-of-the-art super-resolution methods [40], [82]. A summary of the performance of super-resolution augmented face recognition techniques employing LR face inputs from two popular face datasets (AR [46] and YouTube Faces (YTF) [78]) is also presented, and illustrates the performance gap between recognition of faces captured in a controlled environment and recognition of faces captured in the wild. In Section III-C, we propose a center regularization approach for LR face recognition. The goal of this work is to learn a common feature space that locates LR and HR face images of the same subject as close to one another as possible in feature space, without generating thousands of pairs for training. By conducting comprehensive evaluations on two different datasets, we are able to show the performance gap between open-set and closed-set face identification when employing LR inputs. In Section III-D, we first summarize the discoveries in our previous work [37] based on a real surveillance dataset named VBOLO. Different deep architectures are designed and evaluated for LR surveillance face images as a baseline result and are improved by introducing fully convolutional spatial pyramid pooling (SPP). To generalize the ability of

the performance on larger data scale, we further study face re-identification on several surveillance and LR datasets from videos or images collected online. In addition, we present a novel DCGAN-based pre-training approach to further boost performance. In Section IV, concluding remarks are provided.

II. RELATED WORK AND DATASETS

In this section, we present related works on face recognition for LR images and face re-identification, and describe some important datasets that can be used in this area of research.

A. Low-Resolution Face Recognition

The LRFR task is a subset of the general face recognition problem. One of the most useful application scenarios for this task is recognition using surveillance imagery. In this scenario, faces are captured from cameras with a large standoff, positioned above head height, and sometimes under challenging lighting conditions. Even if the faces could be detected by a specially trained face detector, it is hard to construct a robust feature representation due to the lack of information embedded in the image itself. Although HR face recognition systems have achieved nearly perfect performance on several datasets, LRFR remains a challenging problem. Approaches to this problem in the literature follow two main themes. Some techniques employ super-resolution (SR) and deblurring techniques to increase the input LR face size to a point where a high-resolution (HR) face recognition technique may work well. Another popular approach is to learn a unified feature space for LR and HR face images, within which feature vector distance plays its typical role as a matching score.

SR techniques are widely used to turn LR images into better quality images. Face SR approaches are thus an intuitive way to recover LR face images for recognition. Hennings-Yeomans et al. [19] included the prior extracted face features as measures of fit of the SR result, and performs SR from both reconstruction and recognition perspectives. Yang et al. [82] proposed a joint dictionary training method for general SR that employs both LR and HR image patches. Zou et al. [93] designed a linear regression model with two elements (new data and discriminative constraints) to learn the mapping. Jiang et al. [23] proposed a coarse-to-fine face SR approach via a multi-layer locality-constrained iterative neighbor embedding. Kolouri et al. [31] introduced a single frame SR technique that uses a transport-based formulation of this problem. Wang et al. [73] deployed deep learning pre-training with a carefully selected loss function to achieve SR for matching between LR and HR face images, and achieves state-of-the-art performance on a new surveillance dataset.

Another approach to solving the LR face recognition problem is to seek a unified feature space that preserves proximity between faces of different resolutions. Li et al. [35] proposed a coupled mapping method that projects face images with different resolutions into a unified feature space. Biswas et al. [3] simultaneously embedded the LR and HR faces in a common space such that the distances between them in the transformed space approximates the distance between between two HR face images of the same subject. Ren et al. [58]

used a coupled mapping strategy with both HR and LR counterparts for learning the projection directions as well as exploiting discriminant information. Shekhar et al. [61] proposed robust dictionary learning for LR face recognition that shares common sparse codes. The technique of Qiu et al. [54] learned a domain adaptive dictionary to handle the matching of two faces captured in source and target domains. Li et al. [36] proposed several shallow network structures to learn a latent space between LR and HR images, and evaluated the proposed methods on a new surveillance dataset. There are also some other methods such as [56] and [20], which explore more robust features directly to improve recognition rate in blurry and degraded face images. In addition, methods like those proposed in [17], [49] and [50] work to restore the upsampled LR images using deblurring techniques.

In the last three years, novel face recognition methods based on deep learning for low-quality face images have been developed. The most relevant approaches are as follows. In [74], partially coupled networks are proposed for unsupervised super-resolution pre-training. The classifier is obtained by fine-tuning on a different dataset for specific domain simultaneous super-resolution and recognition. In [25], [26], an attention model that shifts the network's attention during training by blurring the images with various percentage of blurriness is presented for gender recognition. In [47], three obfuscation techniques are proposed to restore face images that have been degraded by *mosaicing* (pixelation) and blurring processes. In [4], a multi-task deep model is proposed to simultaneously learn face super-resolution and facial landmark localization. The face super-resolution subnet is trained using a generative adversarial network (GAN) [12], [33] (see a comparison of different versions of GANs in the context of face super-resolution in [8]). In [22], inspired by the traditional wavelet that can depict the contextual and textural information of an image at different levels, a deep architecture is proposed. In [7], a network that contains a coarse super-resolution network to recover a coarse HR image is presented. It is the first deep face super-resolution network utilizing facial geometry prior to end-to-end training and testing. In [10], a network for deblurring facial images using a Resnet-based non-max-pooling architecture is proposed. In [85], a face hallucination method based on an upsampling network and a discriminative network is proposed. The approach includes feature maps with additional facial attribute information. In [62], global semantic priors of the faces are exploited in order to restore blurred face images. In [44], a new branch network that can be appended to a trunk network to match a different resolution of probe images to the gallery images was designed.

In all these methods, we see that automatic face recognition is far from perfect when tackling more challenging images of faces taken in unconstrained environments, *e.g.* surveillance, forensics, etc. Most of the approaches mentioned above are evaluated on low-quality versions of constrained face datasets like Multi-PIE, FERET or FRGC. The images are created by directly downsampling or blurring original images. However, the LR face recognition problem becomes a challenge when faces captured in an unconstrained environment. In such cases, the LR face recognition problem needs to be explored more

extensively.

B. Face Re-Identification

Here, we briefly introduce person re-identification (ReID) and our motivation for face re-identification. A typical end-to-end ReID system generally involves the following steps: person detection, preprocessing, feature extraction, and matching. It is widely used for surveillance purposes, modeled as given a pair of images. The objective is to find whether the images are from the same person or not. Often, the pictures are captured from different cameras in a surveillance network. Traditional approaches to the ReID problem focus on two main components: feature extraction and similarity computation for matching. We can divide most of the existing methods into two categories: methods employing deep learning, and those without. For the non-deep learning methods, research have proposed and used different handcrafted features such as symmetry-driven accumulation of local features (SDALF) [2], color histograms [80], [89], color names [83], local binary patterns [28], aggregation of patch summary features [90], metric learning approaches [80], [28], [30], [38], [45], [39], [80], [24], [6], [88] and various combinations of these. Several deep learning approaches use “Siamese” deep convolutional neural networks [84], [70], [71], [41] for feature extraction and metric learning at the same time providing novel end-to-end solutions. Ahmed et al. [1] proposed a new deep learning framework based on the idea from Yi et al. [84], where two novel layers are employed for computing cross-input neighborhood differences by integrating local relationships based on mid-level features. They additionally showed that the features acquired from the head and neck could be an important clue for person re-identification. Wu et al. [79] improved performance based on Ahmed’s idea by using a deeper architecture and a new optimization method. Other deep network structures such as [69] and [65] have been designed which also effectively solved the ReID problem on older ReID datasets. Qui et al. [54] attempted to perform facial ReID by using domain adaptation methods to reconcile different facial poses; however, their experiments were performed on the Multi-PIE [15] dataset, in which face images have controlled poses and illuminations. Although an increasing number of surveillance cameras have been deployed in public areas, the quality of the video frames is usually low and people captured in the frame are in an uncontrolled pose and illumination condition. Thus, general person re-identification can be a challenging task. As [5] shows, body and gait might play a role in recognizing the target in LR video frames, however, obscuring the target produced a dramatic drop in human-level recognition performance. Also, the face-mask-out experiment in [36] also demonstrates that the face could be an indispensable part of identity recognition. These works help to motivate the LRFR problem as a component of the re-identification problem.

C. Datasets

There are some good surveillance datasets and also some large scale unconstrained face datasets that contain natural LR faces suitable for training and testing LRFR systems.

Most of the LR face images used for research are generated by downsampling a standard face recognition dataset that is collected in a controlled environment. We select the AR dataset to research the LRFR task under uncontrolled scenarios and other unconstrained LR face datasets for more exploration. It is used to illustrate the different data distribution between LR face images artificially generated from high quality controlled face images and those directly collected in unconstrained scenarios such as surveillance camera networks.

1) *AR Face*: The images in the AR database [46] were taken from 100 subjects (50 women and 50 men) with different facial expressions, illumination conditions, and occlusions by sunglasses or scarves under strictly controlled conditions. In our work, these images are used to estimate the baseline performance of face recognition methods on aligned face images in controlled environments.

2) *MegaFace Challenge 2 LR subset*: The MegaFace challenge 2 (MF2) training dataset [48] is the largest (in the number of identities) publicly available facial recognition dataset, with 4.7 million face images and over 672,000 identities. The MF2 dataset is obtained by running the Dlib [29] face detector on images from Flickr [68], yielding 40 million unlabeled faces across 130,154 distinct Flickr accounts. Automatic identity labeling is performed using a clustering algorithm. We performed a subset selection from the MegaFace Challenge 2 training set with tight bounding boxes to generate a LR subset of this dataset. Faces smaller than 50x50 pixels are gathered for each identity, and then we eliminated identities with fewer than five images available. This subset selection approach produced 6,700 identities and 85,344 face images in total. The extraction process does yield some non-face images, as does the original dataset processing. No further data cleaning is conducted on this subset.

3) *YouTubeFaces*: The YouTubeFaces Database [78] is designed for studying the problem of unconstrained face recognition in video. The dataset contains 3,425 videos of 1,595 different people. An average of 2.15 videos is available for each subject.

4) *SCface*: SCface [14] is a database of static images of human faces collected in an uncontrolled indoor environment using five video surveillance cameras of various qualities. The database contains 4,160 static images (in the visible and infrared spectra) of 130 subjects. We choose the HR and LR visible face subset for training and testing.

5) *UCCSface*: The UnConstrained College Students (UCCS) dataset [59] contains HR images captured from an 18-megapixel camera at the University of Colorado at Colorado Springs, capturing people walking on a campus sidewalk from a standoff of 100 to 150 meters, at one frame per second. The dataset consists of more than 70,000 hand-cropped face regions. An identity is manually assigned to many of the faces. We use the database subset that has assigned identities (180 identities total). Although the data are captured by high-definition cameras, the face regions are tiny due to the large standoff and contain a lot of noise and blurriness.

6) *VBOLO face*: This dataset [36], [37] (formerly known as EBOLO) was collected in several sessions at various

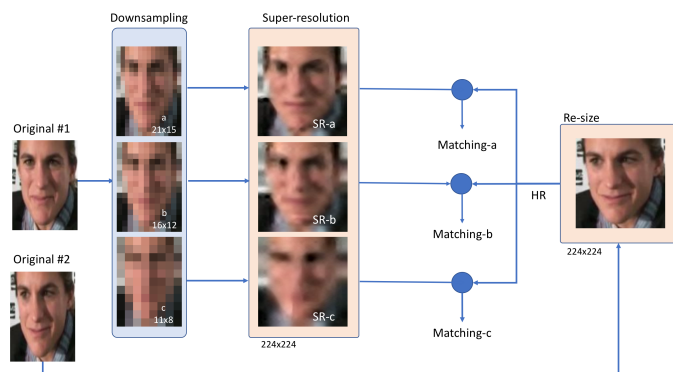


Fig. 1: Super-resolution experiment: two images, #01 (for LR purposes) and #02 (for HR purposes), are matched using five different super-resolution algorithms for two different LR sizes.

checkpoints within public transportation facilities such as tunnels, bridges, and hallways. These capture environments include different camera mount heights and depression angles, illuminations, backgrounds, resolutions, pedestrian poses, and distractors. This dataset provides a good scenario for the facial ReID problem. This dataset uses a small set of known individuals (“actors”) who move in and out of the surveillance cameras’ fields of view, together with unknown persons (“distractors”). The actors change clothing randomly between each appearance in a camera’s field of view. Compared to a typical body-based ReID dataset, which has only a few images for each subject, the VBOLO dataset has a large number of annotations for each subject from consecutive video frames, which mimic a real scenario for surveillance tracking and detection. This is significantly challenging for matching, because *i)* faces change size significantly and exhibit significant pose variations; and *ii)* the two cameras supplying the probe and gallery images may have different resolutions and points of view. In VBOLO, we employ videos captured at two distinct locations (denoted as Station 1 and Station 2) in this research. Each of the collections has nine actors with nine appearances each.

III. METHODS AND EXPERIMENTS

In this Section, we present four groups of methods and experiments on LRFR: *A)* super-resolution techniques, *B)* comparison between virtual and real LR, *C)* face identification and *D)* face re-identification.

A. Super-resolution Techniques

1) Description: In order to explore the gap between the constrained and unconstrained LR face recognition performance, we designed a small super-resolution (SR) experiment with the AR [46] and YouTube Faces (YTF) [78] datasets. In this experiment, the idea is to evaluate the matching performance of two face images: a LR image and a HR image.

In the data selection process, two images or video frames are selected from the datasets for each subject. As illustrated in

Fig. 1, ‘Original #1’ and ‘Original #2’ are used for LR and HR purposes respectively. The first one is downsampled to each of three LR sizes: (a) 21×15 , (b) 16×12 , and (c) 11×8 pixels. The second one is upsampled to the matcher’s required input size of 224×224 pixels. In both cases, bicubic interpolation is used [13]. The images or frames are selected as follows.

- **AR dataset:** original LR and HR face images are taken from face images #01 and #14 respectively of each subject. Since the AR dataset has 100 subjects, we obtained 100 LR–HR pairs of face images for this experiment.

- **YouTube Faces dataset:** original LR and HR face images are chosen as the earliest and latest frame containing a frontal face image in the latest video clip available for each subject. The detection of frontal faces is performed by applying Zhu’s approach [92] to each frame and selecting face images where Zhu’s method returned a pose of 0. The dataset has 1,595 subjects; however, only 1,463 subjects have two suitable frames available in the last video clip. Thus, we have 1,463 LR–HR pairs of face images available.

After the selection of images and frames and scaling to target HR and LR sizes is performed, each of the LR images from that process is then upscaled to the same size of the HR image (224×224 pixels) by each of the following methods: (a) bicubic interpolation [13], (b) SCN [40], (c) sparse representation super-resolution (ScSR) [82], (d) LapSRN [32], and (e) SRGAN [34]. In addition, we included in the experiments the ‘Direct method’, where original images #1 and #2 are compared without downsampling using bicubic interpolation only to achieve the size of 224×224 pixels required by the matcher. This experimental design is depicted in Fig. 1.

2) Experiments and Results: The experimental approach is presented in Fig. 1. The VGG-face trained network in [51] is used to produce a feature vector for the HR image and all fifteen of the SR images. Matching scores are obtained for each match and nonmatch pair involving one HR and one of the fifteen different upscaled SR images. The cosine distance is used as a match score.

The cumulative match characteristic curve for the AR and YouTube Faces datasets can be seen in Figs. 2 and 3 respectively. As might be expected, the performance decreases with decreasing resolution: the Direct method (with no downsampling) achieves better results than those obtained by LR images. Moreover, 21×15 LR-images obtain better performance than 16×12 LR-images, and these ones better than 11×8 LR-images.

In order to show the degradation of performance with low-quality images, we conducted another experiment. Since the original faces in YouTube video frames might be of low quality, we evaluate the performance not only in the matching of all 1,463 pairs, but also in the 500 pairs with the highest quality and in the 500 pairs with lowest quality. We call these two new subsets HQ-YT and LQ-YT respectively. For the measurement of quality of the original face images, we use a score based on the ratio between the high-frequency coefficients and the low-frequency coefficients of the wavelet transform of the image [52]. Low score values indicate low quality. We evaluate the performance at rank 10 when changing the downsampling target resolution in HQ-YT and LQ-YT.

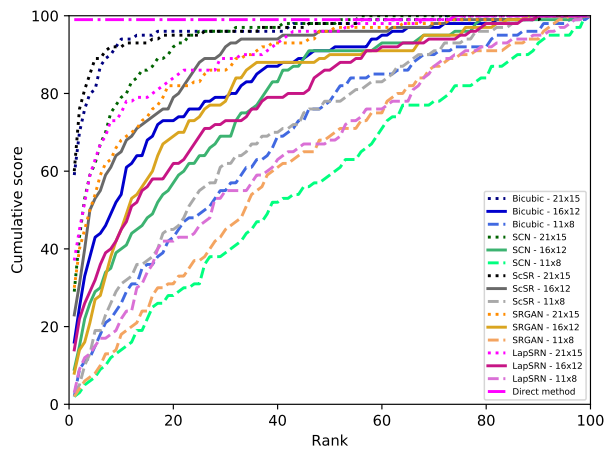


Fig. 2: Cumulative match characteristic for AR.

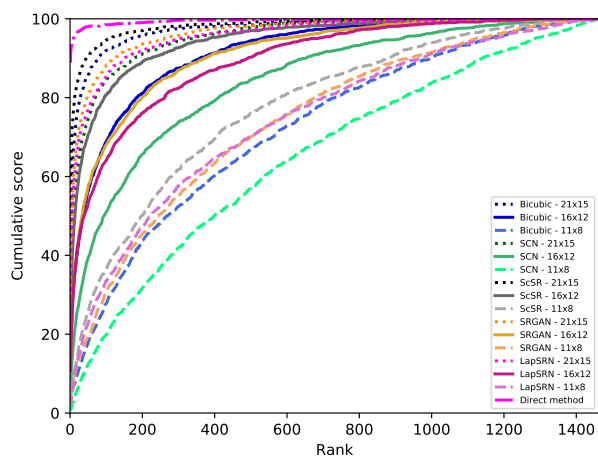


Fig. 3: Cumulative match characteristic for YouTube Faces.

This can be seen in Table I, where high-quality images yield a significantly better performance than the low-quality ones.

The results of Figs. 2 and 3 demonstrate that, for these datasets and these algorithm implementations and at the same resolution, sparse representation super-resolution (ScSR) and bicubic interpolation consistently outperform deep learning super-resolution methods. The poor performance of the deep learning method is due to the introduction of artifacts in severely degraded images as we can see in the example of Fig. 1.

B. Comparison between FR on virtual and real LR images

1) *Description:* In order to explore the gap between virtual (synthetic) LR and real LR, we designed an experiment with the YouTube Faces and SCface datasets to evaluate the gap in performance by matching a HR image and a LR image scaled up with either bicubic or SRGAN between a synthetic

TABLE I: Rank-10 correct match score for subsets HQ-YT and LQ-YT

Subset		(a)	(b)	(c)
		21×15	16×12	11×8
HQ-YT:	Bicubic	80.8%	53.2%	15.6%
	SCN	69.6%	35.2%	9.6%
	ScSR	87.2%	66.4%	26.2%
	LapSRN	78.4%	56.4%	20.4%
	SRGAN	75.4%	53.4%	16.2%
LQ-YT:	Bicubic	81.4%	45.8%	10.8%
	SCN	64.4%	27.2%	7.2%
	ScSR	87.2%	61.0%	14.6%
	LapSRN	64.7%	38.9%	9.8%
	SRGAN	70.9%	40.5%	12.4%

LR dataset (computed from YouTube Faces as explained in Section III-A) and a real LR dataset (SCface). The upscaling methods used are a traditional method (bicubic) and a deep learning method (SRGAN).

• **YouTube Faces dataset:** The same image selection as the super-resolution experiment is used. Then the LR images are downsampled to 16×12 using bicubic interpolation.

• **SCface:** For each of the 130 subjects in the dataset, the HR mugshot and the LR taken by camera 3 at a standoff of 4.20 meters are chosen as gallery images and probe images. The images are resized to 16×12 as well. Then SRGAN and bicubic interpolation are employed to output face images of 224×224 pixels.

2) *Experiments and Results:* The experimental approach is similar to the super-resolution experiment. We use a trained VGG-face network [51] to produce feature vectors for the HR and LR images for both datasets and scaling methods. The cosine distance is used as a score. The cumulative match characteristic curve for 100 random subjects from Youtube Faces and SCface and both methods can be seen in 4.

The decision to use 100 random subjects is taken to make the comparison fairer. If all the subject are used in order to get a 100% performance in rank-1 YouTube Faces would need to get 1463 subjects correctly but SCface only 130 subjects.

Fig. 4 shows that the performance of the super-resolution approaches on virtual LR is consistently much better than the real LR. Hence, in order to characterize system performance on real LR images, real LR face images should be used instead of synthetic LR face images obtained by simply downsampling the HR face images.

C. Low-Resolution Face Identification

In this section, we focus on LR face identification. We first focus on cross-resolution face identification which applies when the enrolled face images are mostly collected in controlled scenarios with HR and LR faces are captured with surveillance cameras with an uncontrolled pose and lighting conditions. This is a challenging recognition task which relies strongly on a good resolution invariant representation.

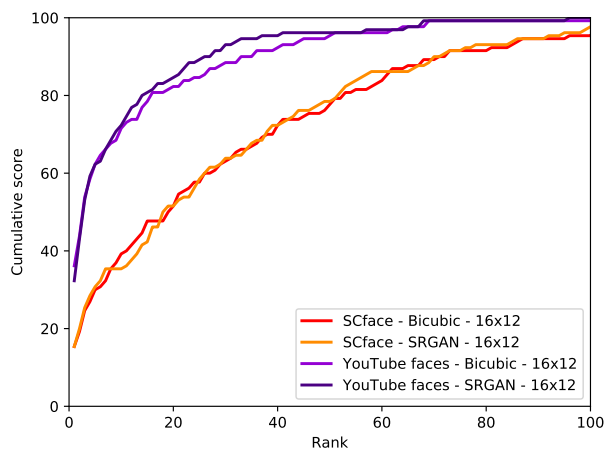


Fig. 4: Cumulative match characteristic plot for virtual and real LR experiment

1) *Description*: Most existing approaches to cross resolution matching would learn a unified space in which LR and HR faces are represented. This requires a carefully designed face pairing mining strategy during the training process which is both time-consuming and performance-sensitive. In order to fully explore the intrinsic connection between the HR and LR domains, we decide to involve the HR and LR face images equally, expecting to learn a common feature space that is able to cluster LR and HR faces within the same subject as well as maintaining low inter-class proximity despite the difference in resolution. We propose an approach based on [77], which employs a novel regularization term to further force the features from the same subject to cluster, which yields better discriminative features. To further stabilize the training process and reduce the overfitting of the model on smaller datasets, we also employ the L2 regularization term. The loss function is:

$$L = - \sum_{i=1}^m \log \frac{e^{W_{yi}^T x_i + b_{yi}}}{\sum_{j=1}^n e^{W_j^T x_i + b_j}} + \frac{2}{\lambda} \sum_{i=1}^m \|x_i - c_{yi}\|_2^2 + \eta \|W\|_2^2$$

where x_i represents batches of face images from different resolutions, and c_{yi} is the center of each class which updates while training.

2) *Experiments and Results*: In most of the previous studies, because of the lack of LR images collected from the wild, HR datasets like Multi-PIE [15] and FERET [53] are used for this task with downsampling of the HR face images to create LR counterparts. Due to the difference between synthetic LR images created by this technique and real surveillance images, the method has limitations when applied to the real world. Two experiment protocols are defined as follows:

- *Closed-set face identification* is typically treated as a classification problem, and the label is predicted directly from the classification layer when a deep learning method is employed. However, when the subjects used for evaluation do not show up in the training stage, the

classification architectures tend to be inflexible enough. This subject-disjoint training and testing protocol are more practical in a real-life scenario.

- *Open-set face identification* requires a 1 to N match when an individual is present to the system. However, it also requires that the system correctly reject individuals who are not enrolled in the system database. This is very similar to how most surveillance systems work. An individual who randomly shows up in the scene may or may not be in the face database. In this case, the system must correctly reject the probe if the person is not in the database and correctly identify if the person is in the database.

Evaluations are conducted on two surveillance quality datasets (SCface and UCCSface) which we consider both more challenging and closer to real scenarios. For the network architecture, we employ a shallower network with seven weight layers, described as conv(32)-conv(32)-maxpooling-conv(64)-conv(64)-maxpooling-conv(128)-conv(128)-maxpooling-flatten-dense. The convolutional kernel size is kept as 3. For center loss, we use a dense layer of 512 dimensions, and for large margin softmax loss [43], additive margin softmax [72] and L2-constrained softmax loss [57], we use a dimension of 256 for the dense layer. For center loss, we set alpha to 0.9. For L2-constrained softmax loss, alpha is set to 5. For large margin softmax loss we set the margin to 1, and for additive margin loss, we set the margin to 0.35.

- **SCface dataset**: Each of the 130 subjects in the subset of SCface dataset that we use has one mugshot HR face image taken by the high-definition camera and 15 LR face images that are captured by five visible light cameras placed at three different standoff distances (1m, 2.6m, and 4.2m). This yields 2080 faces in total. We conduct two sets of experiments with two different protocols as defined in [81]. We split the training set and testing set into 80 and 50 subjects separately in a subject-disjoint fashion. For Experiment 1, we use the HR images as gallery images, and those captured at the three standoffs as probe images. In Experiment 2 we chose face images from 1m standoff as gallery images and images from 2.6m and 4.2m standoff as probe images. The other settings are identical to experiment 1. All the HR and LR images are resized to 64x64 for presentation to the network for training and testing. We conduct the matching using cosine distance of the matching score and the rank-1 rate is reported in II. We achieve nearly five percent improvement in rank-1 rate in Experiment 1 and 9 percent in Experiment 2 compared to the state-of-the-art under the same protocol. The feature generated by our model is more robust when the gallery and probe images have a large resolution level difference. The performance of other methods drops rapidly when the size of the probe image changes from 1m standoff to 4.2m standoff while employing feature extracted from our model.

- **UCCSface dataset**: UCCSface is another dataset designed to be close to a real surveillance setting. We follow the experiment setting provided in [73] and [59] and study both closed-set and open-set scenarios. For closed-set evaluation, 180 subjects are used and for open set evaluation, we compare our result with the performance reported in [59] with the

TABLE II: Experiment 1: Rank-1 rate on SCface with HD and three standoff distances

Method	HD-1m	HD-2.6m	HD-4.7m
SCface [14]	6.18%	6.18%	1.82%
CLPM [35]	3.08%	4.32%	3.46%
SSR [82]	18.09%	13.2%	7.04%
CSCDN [75]	18.97%	13.58%	6.99%
CCA [76]	20.69%	14.85%	9.79%
DCA [16]	25.53%	18.44%	12.19%
C-RSDA [11]	18.46%	18.08%	15.77%
Centerloss [77](ours)	31.71%	20.80%	20.40%
LMsoftmax [43](ours)	18.00%	16.00%	14.00%
AMsoftmax [72](ours)	18.4%	20.8%	14.80%
L2softmax [57](ours)	16.8%	18.8%	9.2%

TABLE III: Experiment 2: Rank-1 rate on SCface with 1.0m and 2.6m standoff distances

Method	1.0m-2.6m
CLPM [35]	29.12%
SDA [91]	40.08%
CMFA [64]	39.56%
Coupled mapping method [63]	43.24%
LMCM [87]	60.40%
Centerloss [77](ours)	69.60%
LMSsoftmax [43](ours)	40.4%
AMSoftmax [72](ours)	46.8%
L2softmax [57](ours)	42.8%

defined openness at 14.11 percent. When looking at the result of the closed-set evaluation, our method beats the UCCS baseline by nearly 20 percent on rank-1 accuracy and also outperforms the DNN method in [73] by nearly 35 percent on rank-1 rate under the same training and evaluation protocol. For open-set evaluation, we achieve 73.6 percent accuracy when compared to the UCCS face baseline result.

TABLE IV: Rank-1 rate on UCCSface dataset

Resolutions	Method	Rank-1 rate
Original vs Original:	UCCSface	78.00%
	Centerloss [77](ours)	95.10%
	LMSsoftmax [43](ours)	65.8%
	AMSoftmax [72](ours)	60.6%
	L2Softmax [57](ours)	86.50%
80×80 vs. 16×16: (HR vs. LR)	DNN	59.03%
	Centerloss [77](ours)	93.40%
	LMSsoftmax [43](ours)	64.9%
	AMsoftmax [72](ours)	58.6%
	L2Softmax [57](ours)	85%

D. Low-resolution face re-identification

In this section, we explore LR face re-identification and evaluate it on several datasets captured in an unconstrained environment. We employ the VBOLO dataset for an in-depth study and the SCface, UCCSface, and MegaFace challenge 2 LR subset for other topical explorations.

1) Experiments and Results:

a) Actor-Disjoint Experiment with selected deep networks: We explore the literature and are inspired by recent state-of-the-art patch matching approaches that may exhibit robustness to small misalignments in automatic or manual annotations and robustness to different effective resolutions. We also expect the network to handle matching of faces captured by the same or different cameras with different subject standoffs. We exploit four state-of-the-art face matching approaches with basic DNN architectures, and boost them with fully convolutional structures to reduce overfitting on our dataset. Further, to let the network better accommodate resolution changes, we employ a spatial pyramid pooling (SPP) layer, hoping to learn discriminative features and the mapping between a different size of LR faces captured in the surveillance cameras.

b) Matching protocols: We make two different matching protocols for this dataset. The first protocol is designed to match people in the same camera between different appearances (usually called “single camera ReID”). This experiment aims to test face identification performance on people from different appearances in the video while clothes are different in the appearances. The second protocol matches images of people acquired from different camera locations. This protocol aims to evaluate the comprehensive performance of the ReID model which includes single-camera and multi-camera person ReID at the same time. In order to increase the matching complexity, we add distractors to the protocol to obtain more non-match pairs.

We train and test several state-of-the-art deep learning patch-matching architectures on our dataset following the two matching protocols. The training and testing sets are disjoint by actor ID in order to mimic the reality that the targets will be unlikely to appear in the training set. Six of the actors are used for training and three for testing. Each set of experiments is conducted five times using the random pair sampling procedure below, and the results are averaged.

c) Training pairing strategy: Creating pairs and sampling the generated pairs in training for matching is a key step of preprocessing. Since we have on average around 200 faces for each person in each appearance, the numbers of positive and negative pairs are highly unbalanced. In addition to pairs created using faces in different appearances, we decide to add face pairs from the same appearance in the training data in order to increase the number of training pairs. We denote the training set as T , and a particular face in T as t_{ijf} , where $i \in 1 \dots n$ is the ID of the actor, $j \in 1 \dots m$ is the appearance number, and f denotes the frame number. We first randomly shuffle all the faces in order to break the temporal continuity of the frames to avoid getting positive face pairs more often from frames close in time to each other. For each of the faces T_{ijf} with fixed i and j , we create a positive pair by randomly

selecting $T_{i,j',f'}$ with $j' \neq j$ and $f' \neq f$. To create an equal number of negative pairs, we randomly select the faces from $T_{i',j,f'}$ with $i' \neq i$ and j' and f' randomly chosen, and pair the selected face with the previous face $T_{i,j,f}$. By exploiting the pairing approaches mentioned above, we are able to gain balanced face pairs from each identity in each appearance.

d) VGG-Face & SVM: We employ the pre-trained VGG-face descriptor model and use it for facial feature extraction. We pair the faces from our dataset first and extract features using the VGG face descriptor [51]. The face images are resized from their original size to 224×224 for input to the network. The Euclidean distance between feature vectors from a pair of faces is assigned with positive and negative binary labels to represent if the pair of faces is from the same person or not. The distances themselves are used to train a linear SVM model for binary prediction. We conduct the experiment with the two matching protocols mentioned above. We chose the hyper-parameter value and obtain the best CV rate for the linear SVM, which yielded a testing AUC of 0.695 as shown in Fig. 6. This result indicates the network successfully identified face features from poor quality faces extracted from surveillance videos. However, since the VGG face descriptor is trained on various HR faces that are sufficiently aligned with facial landmarks, when used for surveillance quality faces that are both LR and hard to align, this baseline result is not outstanding. Also, LR face images need to be upscaled by a large factor before being fed into the deep pre-trained network, which may introduce artifacts and also increase computational complexity.

e) Siamese Network: Siamese classification structures had their first application in face verification in [9]. The Siamese architecture does not require categorical information in training. It tries to learn a feature representation with two identical towers of network layers with shared weights. It has a series of convolutional, activation and max-pooling layers in each tower. Two feature representations coming from each tower are concatenated and fed into fully connected layers which are connected with contrastive loss. Since the two towers are identical in structure and weights, this kind of network aims to map two inputs into an identical target space by training in an end-to-end fashion. We propose a tiny base network with a simple architecture, motivated by [73], which concludes that a deeper network architecture and a large number of filter channels may degrade recognition performance. Our basic network is shown in Table V; it has three convolutional layers followed by max-pooling and a fully connected layer. We use a moderate filter size and channel number in the tiny network. An input size of 32×32 is chosen. We train the network from scratch with a batch size of 8 and the SGD optimizer. The Siamese net converged within 5 epochs with an AUC of 0.861 on Station 1 data and 0.871 on Station 2 data with single camera matching and 0.838 on the data from two stations with both single and cross camera matching.

f) MatchNet: MatchNet [18] is another state-of-the-art patch matching approach that employs a two-tower structure with shared weights similar to the Siamese net. However, instead of feeding two concatenated feature vectors produced by the two towers directly into the decision layer with carefully

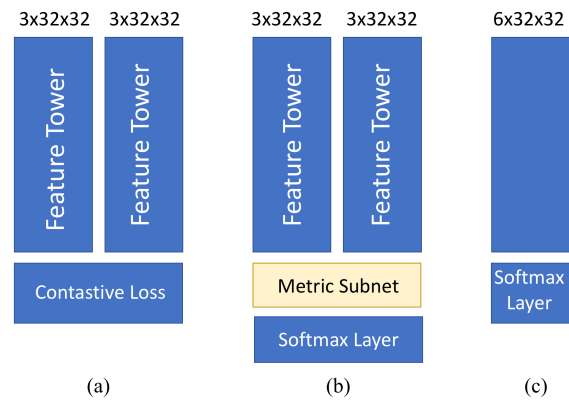
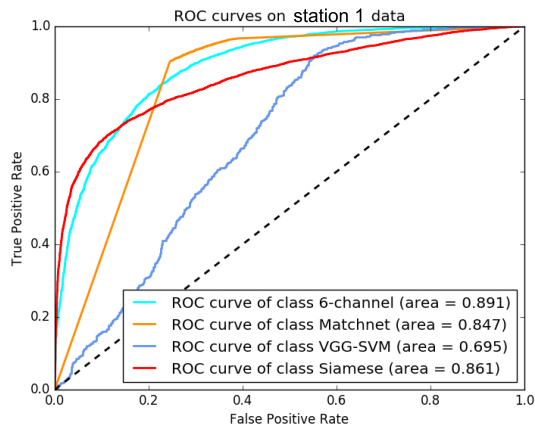


Fig. 5: Overview of the three deep architectures: (a) Siamese net (b) Matchnet (c) 6-channel net

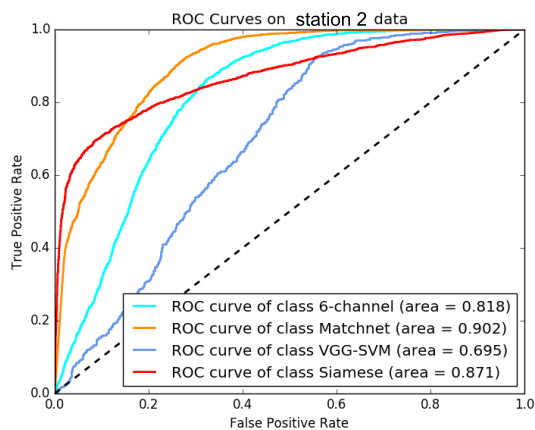
designed loss functions, it uses a series of fully connected layers as a subnet to learn the feature comparison for binary classification using cross-entropy. Compared to the Siamese net, MatchNet has more flexibility in the metric subnet shown in Fig. 5, which takes the paired features and maps them to a unified space that minimizes their distance. However, it converges more slowly since the fully connected layer has many more parameters and higher complexity. A softmax layer and cross-entropy loss are employed during training. We obtain the best result on our basic net using the SGD optimizer. As shown in Fig. 6, it obtains AUC of 0.847 on data from Station 2, 0.902 on data from Station 1 for single camera matching, 0.827 on data from both stations for single and cross-camera matching.

g) Six-Channel Net: Inspired by the two-channel model proposed in [86], we improve it by incorporating three color channels. This approach abandons the two-tower feature by directly embedding the two face images into six channels, fed into the first layer of the network, with hinge loss and a one-bit binary output which is shown in Fig. 5. Compared to the previous two architectures, it has greater flexibility—it has twice as many parameters as the two-tower structures and is able to learn feature maps using six image channels instead of three jointly. However, it converges slowest among the three and L2 regularization is needed for better performance. The best AUC values we achieve using the six-channel net on Station 1 and Station 2 single camera matching are 0.891 and 0.818. It achieves AUC of 0.846, which outperform the Siamese net and MatchNet by 2 percent and 1 percent, respectively.

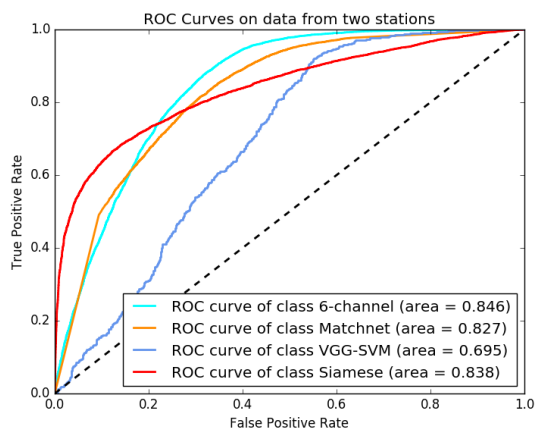
h) Fully convolutional structure and SPP pooling: In this section, we try to improve the three architectures mentioned above (the Siamese net, MatchNet, and the 6-channel net with a fully convolutional structure and a Spatial Pyramid Pooling (SPP) layer [27]). A fully convolutional CNN (FCN) is one where all the learnable layers are convolutional. A convolutional layer has fewer parameters than a fully connected layer, which would potentially reduce overfitting on a small dataset but keep more spatial information in the features. We replace the fully connected layer with a convolutional layer in the



(a)



(b)



(c)

Fig. 6: Face matching with three matching basic net (first line in Table V) for three data subsets. (a): Station 1 (b): Station 2 (c) Station 1 and Station 2.

TABLE V: Model selection for a convolution kernel size of 3×3 . Testing accuracy is shown with different network layouts

Parameters	Method	Accuracy
c(32)-m-c(32)-m-c(64)-fc(64)	6-channel	77.1%
	MatchNet	77.3%
	Siamese	78.3%
c(8)-m-c(16)-m-c(32)-c(32)-m	6-channel	76.7%
	MatchNet	78.1%
	Siamese	77.8%
c(8)-m-c(16)-c(16)m-c(32)-c(32)-m	6-channel	77.5%
	MatchNet	79.3%
	Siamese	79.8%
c(8)-c(8)-m-c(16)-c(16)m-c(32)-c(32)-m	6-channel	76.8%
	MatchNet	75.9%
	Siamese	78.5%

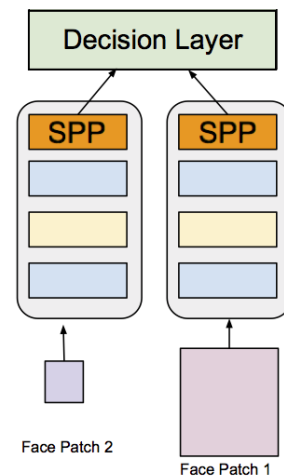


Fig. 7: Architecture of Siamese network with SPP layer.

previous three models, test several hyperparameter settings to adjust layer numbers and filter numbers, and chose the best settings based on the testing accuracy observed. Performance comparison is summarized in Table V. Fig. 8 demonstrates that the FCN architecture effectively improves the performance (as measured by the AUC) over the three basic network architectures (Siamese, Matchnet, 6-channel net) by 1 percent, 5 percent and 4 percent roughly each. Further, we are inspired by the study of Zagoruyko et al. [86] who used an SPP layer for patch matching and claimed a remarkable improvement in performance. Compared with [86] who only tested their architecture with SPP using size-identical image pairs, we decide to take advantage of the fully connected architecture of our network, feeding various size of face images into the network. By replacing the last max-pooling layer with an SPP layer before the decision layer, we can further test the assumption from [86] on the VBOLO dataset with our modified architecture. We simplify the problem by setting up three convolutional Siamese networks. Each is responsible for a matching at a specific resolution level. In this case, we have

a) low to low, b) high to high, and c) low to high-resolution matching with three separate but identical networks. These three networks try to learn different metrics and features with face pairs close to their original sizes. We resize faces that are smaller than 32×32 to 16×16 and denote them as LR faces. Those faces bigger than 32×32 are resized to 64×64 and denoted as HR faces. We train and test the three subnets with the SPP layer on the top shown in Fig. 7. 4×4 SPP pooling is applied at the end of each tower. We got a slight (0.1 percent) AUC improvement using the SPP layer together with the fully convolutional architecture. Compared with previous work [36], we achieve a significant improvement in performance on the data from Station 2 by exploiting the unified deep feature and metric learning instead of optimizing feature and metric separately.

i) Improved pretraining approach on larger datasets:

Although we achieve several sets of promising results on the VBOLO dataset, more comprehensive experiments need to be conducted on larger scale public datasets. In addition, we face the challenge that most of the deep architectures struggle to exploit LR images well due to over-fitting and limited intrinsic dimensionality of the input. We execute a project to improve the training process by evaluating on some larger general LR unconstrained face datasets or surveillance quality face datasets. To achieve this goal, we employ the DCGAN introduced in [55] in order to obtain a pre-trained discriminator as an initialization for the feature towers. Using this method has two advantages:

- By optimizing the DCGAN on the LR training set, we can understand how the network perceives the LR images and adjust parameters such as activation function and the number of layers by looking at the intermediate output of the generator as shown in Fig. 9. Visualizations from filters and feature maps are not yet intuitive enough to inform a training strategy.
- The pretrained discriminator could provide an initial weight on general LR faces which can be transferred to other LR face datasets via fine-tuning which can stabilize and accelerate the training process.

The GAN discriminator is trained on the MegaFace challenge 2 LR subset and fine-tuned using the target datasets (VBOLO, SCface and UCCSface). We compare the model trained from scratch and a fully convolutional MatchNet model pre-trained using DCGAN.

We perform the training and testing splitting as follows: For UCCSface, we choose 90 identities for training and 90 for testing. For MegaFace challenge 2 LR subset, we use 2999 identities for training and 6699 identities for testing. For SCface, we select 65 identities for training and 65 for testing. We follow the VBOLO pairing and matching protocol for training and testing. The validation set is a randomly selected 20 percent of the training set. All the experiments are run five times and the error rates are averaged. When starting with a pretrained weight using a larger number of LR faces by applying the DCGAN discriminator, we observe that the validation rate is more stable and tended to be higher compared with training the network from scratch. Better performance

TABLE VI: Average error rate on Matchnet models

Datasets	VBOLO	UCCSface	MegaFace 2 sub	SCface
DCGAN-pretrained	18.8/17.6	14.7/11.8	20.1/19.8	24.3/24.1
Train from scratch	20.7/19.5	18.8/18.6	22.5/21.8	27.4/28.5

is achieved on the same datasets. For VBOLO, UCCSface, MegaFace challenge 2 LR subset and SCface, we achieve 9.2, 21.8, 10.6, 11.3 percent decrease in error rate compared with training the model from scratch.

We also identify some architectural changes needed for a DCGAN model to successfully converge on LR face images compared with on HR images. For higher resolution modeling, Radford et al. [55] suggested that the following steps would result in stable training:

- Replace any pooling layers with strided convolutions (discriminator) and fractional-strided convolutions (generator).
- Replace Tanh activation function with ReLU or Leakyrelu functions.
- Add batch-normalization in both generator and discriminator.

However, we only achieve stable adversarial convergence using the Tanh activation function in both the generator and the discriminator except for the last layer of the discriminator, in which we employ a sigmoid nonlinearity. Batch-normalization does not function usefully for our DCGAN to converge on the MegaFace LR subset and is not applied in our model.

IV. CONCLUSIONS AND COMMENTARY

In this paper, we provide several novel contributions. First, we illustrate the performance gap between LR unconstrained face and LR constrained face recognition when using a state-of-the-art super-resolution algorithm. Secondly, two important application scenarios based on LR face recognition are defined: unconstrained LR face identification in the wild and LR face re-identification. For general LR face identification, we exploit a novel approach to handle the multi-dimensional mismatching due to the quality difference of the face images in probe and gallery. We also design different deep networks solving the person re-identification problem to demonstrate better performance compared to our previous work [36], [37]. We present a novel strategy using DCGAN pre-training to obtain both the learning visualization of the network and improve result on larger scale datasets. We demonstrate the result from the extensive experiments on selected datasets and discover that dimensional mismatching is the most challenging point, especially in low-to-high resolution face identification task. The result demonstrates that the approaches we proposed target different tasks, work efficiently, and yield impressive performance.

REFERENCES

[1] E. Ahmed, M. Jones, and T. K. Marks. An improved deep learning architecture for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3908–3916, 2015.

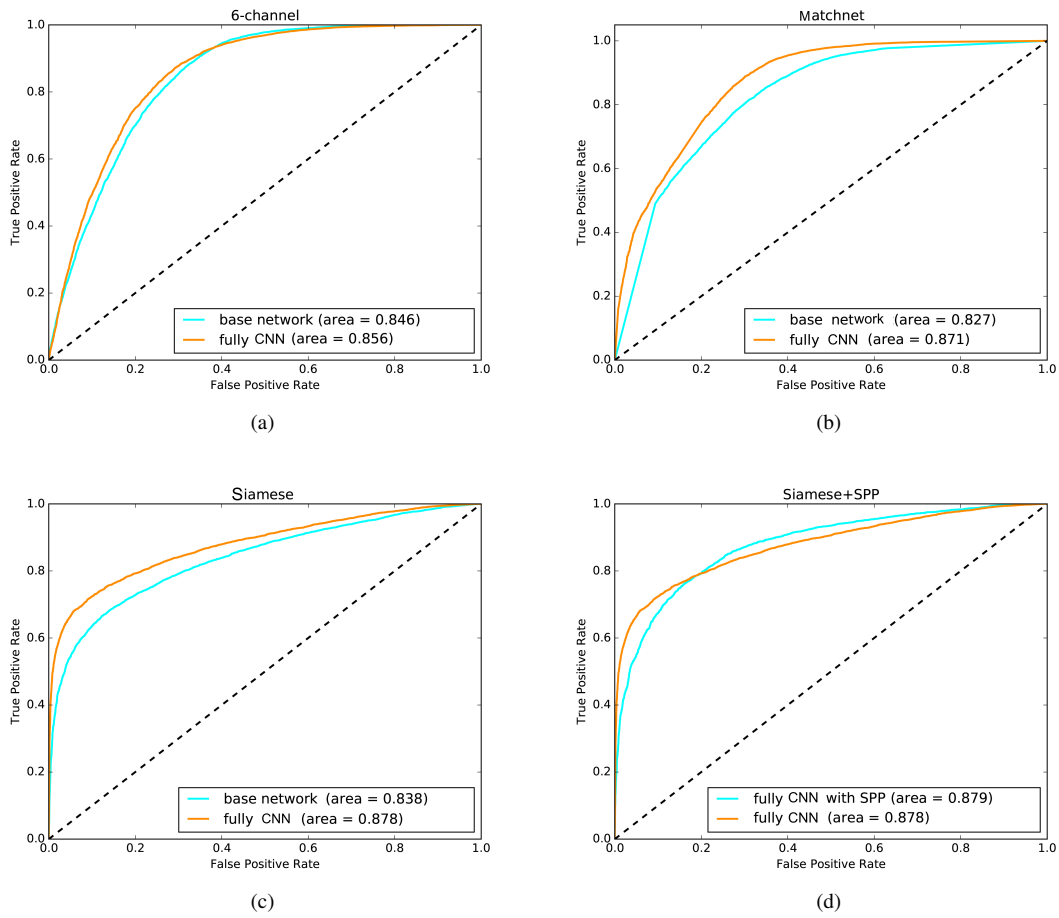


Fig. 8: Face matching with fully convolutional networks on faces from two stations with multi-camera and multi-appearances. (a): 6-channel. (b): Matchnet. (c): Siamese net. (d): Siamese new with SPP layer.



Fig. 9: Generated LR faces from DCGAN generator

- [2] L. Bazzani, M. Cristani, and V. Murino. Symmetry-driven accumulation of local features for human characterization and re-identification. *Comput. Vis. Image Underst.*, 117(2):130–144, 2013.
- [3] S. Biswas, K. W. Bowyer, and P. J. Flynn. Multidimensional scaling for matching low-resolution facial images. In *Proceedings of the IEEE Conference on Biometrics: Theory Applications and Systems (BTAS)*, pages 1–6, 2010.
- [4] A. Bulat and G. Tzimiropoulos. Super-FAN: Integrated facial landmark localization and super-resolution of real-world low resolution faces in arbitrary poses with GANs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [5] A. M. Burton, S. Wilson, M. Cowan, and V. Bruce. Face recognition in poor-quality video: Evidence from security surveillance. *Psychological Science*, 10(3):243–248, 1999.
- [6] D. Chen, Z. Yuan, B. Chen, and N. Zheng. Similarity learning with spatial constraints for person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1268–1277, 2016.
- [7] Y. Chen, Y. Tai, X. Liu, C. Shen, and J. Yang. FSRNet: End-to-end learning face super-resolution with facial priors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [8] Z. Chen and Y. Tong. Face super-resolution through wasserstein GANs. *arXiv*, 2017.
- [9] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, pages 539–546, 2005.
- [10] G. G. Chryso and S. Zafeiriou. Deep face deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2017.
- [11] Y. Chu, T. Ahmad, G. Bebis, and L. Zhao. Low-resolution face recognition with single sample per person. *Signal Processing*, 141:144–157, December 2017.
- [12] A. Creswell, T. White, V. Dumoulin, K. Arulkumaran, B. Sengupta, and A. A. Bharath. Generative adversarial networks: an overview. *IEEE Signal Process Magazine*, 35(1):53–65, Jan 2018.
- [13] R. Gonzalez, R. Woods, and S. Eddins. *Digital Image Processing*. Prentice Hall Press, 2009.
- [14] M. Grgic, K. Delac, and S. Grgic. Seface—surveillance cameras face database. *Multimedia Tools and Applications*, 51(3):863–879, 2011.
- [15] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. *Image and Vision Computing*, 28(5):807–813, 2010.
- [16] M. Haghighat and M. Abdel-Mottaleb. Low resolution face recognition in surveillance systems using discriminant correlation analysis. In *proceedings of 12th IEEE International Conference on Automatic Face & Gesture Recognition*, pages 912–917, 2017.
- [17] Z. Haichao, J. Yang, Y. Zhang, N. M. Nasrabadi, and T. S. Huang. Close the loop: Joint blind image restoration and recognition with sparse representation prior. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 770–777, Nov, 2011.
- [18] X. Han, T. Leung, Y. Jia, R. Sukthankar, and A. C. Berg. Matchnet: Unifying feature and metric learning for patch-based matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3279–3286, 2015.
- [19] P. H. Hennings-Yeomans, S. Baker, and B. V. Kumar. Simultaneous super-resolution and feature extraction for recognition of low-resolution faces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [20] C. Herrmann, D. Willersinn, and J. Beyerer. Low-resolution convolutional neural networks for video face recognition. In *Proceedings of the IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 221–227, 2016.
- [21] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007.
- [22] H. Huang, R. He, Z. Sun, and T. Tan. Wavelet-SRNet: A wavelet-based CNN for multi-scale face super resolution. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1698–1706, Oct 2017.
- [23] J. Jiang, R. Hu, Z. Wang, and Z. Han. Face super-resolution via multilayer locality-constrained iterative neighbor embedding and intermediate dictionary learning. *IEEE Transactions on Image Processing*, 23(10):4220–4231, Oct 2014.
- [24] C. Jose and F. Fleuret. Scalable metric learning via weighted approximate rank component analysis. *arXiv preprint arXiv:1603.00370*, 2016.
- [25] F. Juefei-Xu, E. Verma, P. Goel, A. Cherodan, and M. Savvides. DeepGender: Occlusion and low resolution robust facial gender classification via progressively trained convolutional neural networks with attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 136–145. IEEE, Jun 2016.
- [26] F. Juefei-Xu, E. Verma, and M. Savvides. DeepGender2: A generative approach toward occlusion and low-resolution robust facial gender classification via progressively trained attention shift convolutional neural networks (PTAS-CNN) and deep convolutional generative adversarial networks (DCGAN). In *Deep Learning for Biometrics*, pages 183–218. Springer, 2017.
- [27] H. Kaiming, Z. Xiangyu, R. Shaoqing, and J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. In *Proceedings of the IEEE European Conference on Computer Vision*, 2014.
- [28] S. Khamis, C.-H. Kuo, V. K. Singh, V. D. Shet, and L. S. Davis. Joint learning for attribute-consistent person re-identification. In *Proceedings of the European Conference on Computer Vision*, 2014.
- [29] D. E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10(Jul):1755–1758, 2009.
- [30] M. Koestinger, M. Hirzer, P. Wohlhart, P. M. Roth, and H. Bischof. Large scale metric learning from equivalence constraints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2288–2295, 2012.
- [31] S. Kolouri and G. K. Rohde. Transport-based single frame super resolution of very low resolution face images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4876–4884, 2015.
- [32] W.-S. Lai, J.-B. Huang, N. Ahuja, and M.-H. Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [33] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 105–114. IEEE, jul 2017.
- [34] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. P. Aitken, A. Tejani, J. Totz, Z. Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, page 4, 2017.
- [35] B. Li, H. Chang, S. Shan, and X. Chen. Low-resolution face recognition via coupled locality preserving mappings. *IEEE Signal Processing Letters*, 17(1):20–23, Jan 2010.
- [36] P. Li, J. Brogan, and P. J. Flynn. Toward facial re-identification: Experiments with data from an operational surveillance camera plant. In *Proceedings of the IEEE Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–8, 2016.
- [37] P. Li, M. L. Prieto, P. J. Flynn, and D. Mery. Learning face similarity for re-identification from real surveillance video: A deep metric solution. In *Proceedings of the IEEE International Joint Conference on Biometrics (IJCB)*, Oct 2017.
- [38] W. Li and X. Wang. Locally aligned feature transforms across views. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3594–3601, 2013.
- [39] Z. Li, S. Chang, F. Liang, T. Huang, L. Cao, and J. Smith. Learning locally-adaptive decision functions for person verification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3610–3617, 2013.
- [40] D. Liu, Z. Wang, B. Wen, J. Yang, W. Han, and T. S. Huang. Robust single image super-resolution via deep networks with sparse prior. *IEEE Transactions on Image Processing*, 25(7):3194–3207, 2016.
- [41] H. Liu, J. Feng, M. Qi, J. Jiang, and S. Yan. End-to-end comparative attention networks for person re-identification. *arXiv preprint arXiv:1606.04404*, 2016.
- [42] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song. Sphereface: Deep hypersphere embedding for face recognition. In *Proceedings*

- of the *IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, 2017.
- [43] W. Liu, Y. Wen, Z. Yu, and M. Yang. Large-margin softmax loss for convolutional neural networks. In *In Proceedings of the International Conference on Machine Learning*, pages 507–516, 2016.
- [44] Z. Lu, X. Jiang, and A. C. Kot. Deep coupled resnet for low-resolution face recognition. *IEEE Signal Processing Letters*, 2018.
- [45] N. Martinel, C. Micheloni, and G. L. Foresti. Saliency weighted features for person re-identification. In *Proceedings of the IEEE European Conference on Computer Vision Workshops*, pages 191–208, 2014.
- [46] A. M. Martinez. The ar face database. *CVC technical report*, 1998.
- [47] R. McPherson, R. Shokri, and V. Shmatikov. Defeating image obfuscation with deep learning. *arXiv preprint arXiv:1609.00408*, 2016.
- [48] A. Nech and I. Kemelmacher-Shlizerman. Level playing field for million scale face recognition. *arXiv preprint arXiv:1705.00393*, 2017.
- [49] M. Nishiyama, A. Hadid, H. Takeshima, J. Shotton, T. Kozakaya, and O. Yamaguchi. Facial deblur inference using subspace analysis for recognition of blurred faces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(4):838–845, April 2011.
- [50] J. Pan, Z. Hu, Z. Su, and M.-H. Yang. Deblurring face images with exemplars. In *Proceedings of the European Conference on Computer Vision*, pages 47–62, 2014.
- [51] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *Proceedings of the British Machine Vision Conference*, Swansea, 2015.
- [52] S. Pertuz, D. Puig, and M. Garcia. Analysis of focus measure operators for shape-from-focus. *Pattern Recognition*, 46(5):1415–1432, 2013.
- [53] P. J. Phillips, H. Wechsler, J. Huang, and P. J. Rauss. The feret database and evaluation procedure for face-recognition algorithms. *Image and Vision Computing*, 16(5):295–306, 1998.
- [54] Q. Qiu, J. Ni, and R. Chellappa. Dictionary-based domain adaptation methods for the re-identification of faces. In *Person Re-Identification*, pages 269–285. Springer, 2014.
- [55] A. Radford, L. Metz, and S. Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [56] E. Rahtu, J. Heikkilä, V. Ojansivu, and T. Ahonen. Local phase quantization for blur-insensitive image analysis. *Image and Vision Computing*, 30(8):501–512, 2012.
- [57] R. Ranjan, C. D. Castillo, and R. Chellappa. L2-constrained softmax loss for discriminative face verification. *arXiv preprint arXiv:1703.09507*, 2017.
- [58] C.-X. Ren, D.-Q. Dai, and H. Yan. Coupled kernel embedding for low-resolution face image recognition. *IEEE Transactions on Image Processing*, 21(8):3770–3783, 2012.
- [59] A. Sapkota and T. E. Boult. Large scale unconstrained open set face database. In *Proceedings of the IEEE Conference on Biometrics: Theory, Applications and Systems (BTAS)*, pages 1–8, 2013.
- [60] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 815–823, 2015.
- [61] S. Shekhar, V. M. Patel, and R. Chellappa. Synthesis-based recognition of low resolution faces. In *Proceedings of the International Joint Conference on Biometrics (IJCB)*, pages 1–6, Oct 2011.
- [62] Z. Shen, W. Lai, T. Xu, and J. Kautz. Deep semantic face deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. openaccess.thecvf.com, 2018.
- [63] J. Shi and C. Qi. From local geometry to global structure: Learning latent subspace for low-resolution face image recognition. *IEEE Signal Processing Letters*, 22(5):554–558, 2015.
- [64] S. Siena, V. N. Boddeti, and B. V. Kumar. Coupled marginal fisher analysis for low-resolution face recognition. In *proceeding of European Conference on Computer Vision*, pages 240–249. Springer, 2012.
- [65] A. Subramaniam, M. Chatterjee, and A. Mittal. Deep neural networks with inexact matching for person re-identification. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 2667–2675, 2016.
- [66] Y. Sun, Y. Chen, X. Wang, and X. Tang. Deep learning face advancement by joint identification-verification. In *Proceedings of the Advances in Neural Information Processing Systems*, pages 1988–1996, 2014.
- [67] Y. Sun, X. Wang, and X. Tang. Sparsifying neural network connections for face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4856–4864, 2016.
- [68] B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016.
- [69] E. Ustinova, Y. Ganin, and V. Lempitsky. Multiregion bilinear convolutional neural networks for person re-identification. *arXiv preprint arXiv:1512.05300*, 2015.
- [70] R. R. Varior, M. Haloi, and G. Wang. Gated siamese convolutional neural network architecture for human re-identification. In *Proceedings of the European Conference on Computer Vision*, pages 791–808. Springer, 2016.
- [71] R. R. Varior, B. Shuai, J. Lu, D. Xu, and G. Wang. A siamese long short-term memory architecture for human re-identification. In *Proceedings of the European Conference on Computer Vision*, pages 135–153, 2016.
- [72] F. Wang, J. Cheng, W. Liu, and H. Liu. Additive margin softmax for face verification. *IEEE Signal Processing Letters*, 25(7):926–930, 2018.
- [73] Z. Wang, S. Chang, Y. Yang, D. Liu, and T. S. Huang. Studying very low resolution recognition using deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4792–4800, 2016.
- [74] Z. Wang, S. Chang, Y. Yang, D. Liu, and T. S. Huang. Studying very low resolution recognition using deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4792–4800. IEEE, jun 2016.
- [75] Z. Wang, D. Liu, J. Yang, W. Han, and T. Huang. Deep networks for image super-resolution with sparse prior. In *Proceedings of IEEE International Conference on Computer Vision*, pages 370–378, Dec 2015.
- [76] Z. Wang, W. Yang, and X. Ben. Low-resolution degradation face recognition over long distance based on cca. *Neural Computing and Applications*, 26(7):1645–1652, 2015.
- [77] Y. Wen, K. Zhang, Z. Li, and Y. Qiao. A discriminative feature learning approach for deep face recognition. In *Proceedings of the European Conference on Computer Vision*, pages 499–515, 2016.
- [78] L. Wolf, T. Hassner, and I. Maoz. Face recognition in unconstrained videos with matched background similarity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 529–534, 2011.
- [79] L. Wu, C. Shen, and A. v. d. Hengel. Personnet: Person re-identification with deep convolutional neural networks. *arXiv preprint arXiv:1601.07255*, 2016.
- [80] F. Xiong, M. Gou, O. Camps, and M. Szaier. Person re-identification using kernel-based metric learning methods. In *Proceedings of the European Conference on Computer Vision*, pages 1–16. Springer, 2014.
- [81] F. Yang, W. Yang, R. Gao, and Q. Liao. Discriminative multidimensional scaling for low-resolution face recognition. *IEEE Signal Processing Letters*, 25(3):388–392, 2018.
- [82] J. Yang, J. Wright, T. S. Huang, and Y. Ma. Image super-resolution via sparse representation. *IEEE Transactions on Image Processing*, 19(11):2861–2873, 2010.
- [83] Y. Yang, J. Yang, J. Yan, S. Liao, D. Yi, and S. Z. Li. Salient color names for person re-identification. In *Proceedings of the European Conference on Computer Vision*, pages 536–551. Springer, 2014.
- [84] D. Yi, Z. Lei, and S. Z. Li. Deep metric learning for practical person re-identification. *arXiv preprint arXiv:1407.4979*, 2014.
- [85] X. Yu, B. Fernando, R. Hartley, and F. Porikli. Super-resolving very low-resolution face images with supplementary attributes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 908–917. IEEE, 2018.
- [86] S. Zagoruyko and N. Komodakis. Learning to compare image patches via convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4353–4361, 2015.
- [87] J. Zhang, Z. Guo, X. Li, and Y. Chen. Large margin coupled mapping for low resolution face recognition. In *Pacific Rim International Conference on Artificial Intelligence*, pages 661–672, 2016.
- [88] L. Zhang, T. Xiang, and S. Gong. Learning a discriminative null space for person re-identification. *arXiv preprint arXiv:1603.02139*, 2016.
- [89] Z. Zhang, Y. Chen, and V. Saligrama. A novel visual word co-occurrence model for person re-identification. In *Proceedings of the European Conference on Computer Vision Workshops*, pages 122–133, 2014.
- [90] R. Zhao, W. Ouyang, and X. Wang. Person re-identification by saliency matching. In *Proceedings of the IEEE International Conference on*

- Computer Vision*, pages 2528–2535, 2013.
- [91] C. Zhou, Z. Zhang, D. Yi, Z. Lei, and S. Z. Li. Low-resolution face recognition via simultaneous discriminant analysis. In *Proceedings of the IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–6, 2011.
- [92] X. Zhu and D. Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2879–2886, 2012.
- [93] W. W. Zou, P. C. Yuen, and R. Chellappa. Low-resolution face tracker robust to illumination variations. *IEEE Transactions on Image Processing*, 22(5):1726–1739, 2013.