

FAN: Feature Adaptation Network for Surveillance Face Recognition and Normalization

Xi Yin^{†§}, Ying Tai[‡], Yuge Huang^{†◊}, Xiaoming Liu[†]

[†]Michigan State University, [‡]Youtu Lab, Tencent

[§]Microsoft Cloud and AI, [◊]Zhejiang University

{yinxil, liuxm}@cse.msu.edu, {yingtai, yugehuang}@tencent.com




Figure 1: **Visual results on 4 datasets.** Vertically we show input in row 1 and our results in row 3. For LFW and SCface datasets, we show the ground truth and gallery images in row 2, respectively. For WIDER FACE and QMUL-SurFace datasets which do not have ground truth high-resolution images, we compare with two state-of-the-art SR methods: Bulat et al. [4] and FSRGAN [8] in row 2, respectively.

Abstract

This paper studies face recognition (FR) and normalization in surveillance imagery. Surveillance FR is a challenging problem that has great values in law enforcement. Despite recent progress in conventional FR, less effort has been devoted to surveillance FR. To bridge this gap, we propose a Feature Adaptation Network (FAN) to jointly perform surveillance FR and normalization. Our face normalization mainly acts on the aspect of image resolution, closely related to face super-resolution. However, previous face super-resolution methods require paired training data with pixel-to-pixel correspondence, which is typically unavailable between real low- and high-resolution faces. Our FAN can leverage both paired and unpaired data as we disentangle the features into identity and non-identity components and adapt the distribution of the identity features, which breaks the limit of current face super-resolution methods. We further propose a random scale augmentation scheme to learn resolution robust identity features, with advantages over previous fixed scale augmentation. Extensive experiments on LFW, WIDER FACE, QMUL-SurFace and SCface datasets have demonstrated the superiority of our proposed method compared to the state of the arts on

surveillance face recognition and normalization.

1. Introduction

Surveillance Face Recognition (FR) is a challenge and important problem, yet less studied. The performance on conventional benchmarks such as LFW [19] and IJB-A [25] have been greatly improved by state-of-the-art (SOTA) FR methods [11, 43, 44], which still suffer when applied to surveillance FR. One intuitive approach is to perform Face Super-Resolution (FSR) on surveillance faces to enhance facial details. However, existing FSR methods are problematic to handle surveillance faces, because they usually ignore the *identity* information and require *paired* training data. In fact, preserving identity information is more crucial for surveillance faces than recovering other information, e.g., background, Pose, Illumination, Expression (PIE).

In this work, we study surveillance face recognition and normalization. Specifically, given a surveillance face image, we aim to learn robust identity features for FR. Meanwhile, the features are used to generate a normalized face with enhanced facial details and neutral PIE. Our normalization is performed mainly on the aspect of resolution.

⁰Most of this work was done when Xi visited Youtu Lab. Yuge is an Intern at Youtu Lab during this collaboration.

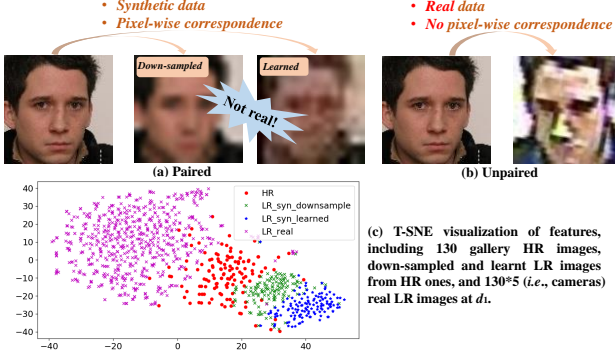


Figure 2: **Paired vs. unpaired data from SCface [13].** Synthetic paired data can be obtained by either down-sampling [8, 62] or via a learned degradation mapping [4].

While sharing the same goal as traditional SR, it differs in removing the pixel-to-pixel correspondence between the original and super-resolved images, as required by traditional SR. Therefore, we term it as face *normalization*. For the same reason, we compare ours to FSR, instead of prior normalization methods operating on pose [42], or expression [63]. To the best of our knowledge, this is the *first* work to study surveillance face normalization.

We propose a novel Feature Adaptation Network (FAN) to jointly perform face recognition and normalization, which has three advantages over conventional FSR. 1) Our joint learning scheme can benefit each other while most FSR methods do not consider the recognition task. 2) Our framework enables training with both paired and unpaired data while conventional SR methods only support paired training. 3) Our method simultaneously improves the resolution and alleviates the background and PIE from real surveillance faces while conventional methods only act on the resolution. Examples in Fig. 1 demonstrate the superiority of FAN over SOTA SR methods.

Our FAN consists of two stages. In the first stage, we adopt disentangled feature learning to learn both identity and non-identity features mainly from high-resolution (HR) images, which are combined as the input to a decoder for pixel-wise face recovering. In the second stage, we propose feature adaptation to further facilitate the feature learning from the low-resolution (LR) images, by approximating the feature distribution between the LR and HR identity encoders. There are two advantages to use FAN for surveillance FR and normalization. First, FAN focuses on learning disentangled identity features from LR images, which is better for FR than extracting features from super-resolved faces [42, 46, 57]. Second, our adaptation is performed in the disentangled identity feature space, which enables training with unpaired data without pixel-to-pixel correspondence. As shown in Fig. 2, the synthetic paired data used in prior works [4, 8, 36, 40, 41, 46, 57, 62] cannot accurately reflect the difference between real LR and HR in-the-wild faces,

which is also observed in [5].

Furthermore, to better handle surveillance faces with unknown and diverse resolutions, we propose a Random Scale Augmentation (RSA) method that enables the network to learn all kinds of scales during training. Prior FSR [8, 55, 62] methods either *artificially* generate the LR images from the HR ones by simple *down-sampling*, or *learn* the degradation mapping via a Convolutional Neural Network (CNN). However, their common drawback is to learn reconstruction under *fixed* scales, which may greatly limit their applications to surveillance faces. In contrast, our RSA efficiently alleviates the constraint on scale variation.

In summary, the contributions of our work include:

- We propose a novel FAN to address surveillance face recognition and normalization, which is suitable for both paired and unpaired data.
- We integrate disentangled feature learning to learn identity and non-identity features, which helps achieve face normalization for visualization, and *identity preserving* for face recognition, simultaneously.
- We propose a random scale augmentation strategy in FAN to learn various scales during training, which addresses the unknown resolutions of surveillance faces.
- We achieve state-of-the-art performances on surveillance datasets: WIDER FACE, QMUL-SurvFace and SCface, both quantitatively and qualitatively.

2. Related Work

Face Recognition Face recognition is a long-standing topic in computer vision. The performance has improved substantially due to the success of CNNs and large training sets [15]. Most previous FR methods are focused on designing better loss functions to learn more discriminative features [11, 43, 44, 51]. For example, Deng et al. [11] proposed ArcFace to introduce a margin in the angular space to make training more challenging and thus learn a more discriminative feature representation. Other methods are proposed to handle one or more specific variations in FR. For example, pose-invariant FR [42, 52, 59] has been widely studied as pose is one major challenge in FR. With recent advance, the performance on conventional benchmarks like LFW [19], CFP [37], and IJB-A [25] are saturating.

However, most previous FR methods fail to achieve satisfactory performance on surveillance FR [9], which is a more challenging task that mainly tackles unconstrained LR faces. To address surveillance FR, one common approach is to learn a unified feature space for LR and HR images [29, 38]. Besides, face SR methods that preserve the identity information are another direction. E.g., Hennings-Yeomans et al. [18] incorporated face features based prior in SR. Wu et al. [46] integrated a FR net after a standard SR net, and jointly learned a deep model for face hallucination and recognition. More recently, Zhang et al. [57] defined a

Table 1: **Comparisons with previous state-of-the-art face super-resolution methods.** Vis., Rec., Dis. and Frontal. indicate visualization, recognition, disentangled and frontalization, respectively. $A \rightarrow B$ refers to performing A first and then followed by B.

Method	FSRNet [8] (CVPR'18)	Bulat et al. [4] (ECCV'18)	S ² R ² [18] (CVPR'08)	Wu et al. [46] (Arxiv'16)	SICNN [57] (ECCV'18)	FAN (Ours)
Deep Model	✓	✓	×	✓	✓	✓
Applications	Vis.	Vis.	Vis. & Rec.	Vis. & Rec.	Vis. & Rec.	Vis. & Rec.
Pipeline	SR	SR	Features & SR → Rec.	SR → Rec.	SR → Rec.	Dis. Features → SR & Frontal. & Rec.
Scale Factors	8	4	2/4	4	8	Random
Identity Preserving	×	×	✓	✓	✓	✓
Scenarios	CelebA [32]/Helen [27] (Easy)	WIDER FACE [50] (Hard)	MPIE [14]/FRGC [34] (Easy)	LFW [19]/YTF [45] (Medium)	LFW [19] (Medium)	SCface [13]/QMUL-SurFace [9]/ WIDER FACE [50] (Hard)

super-identity loss to measure the identity difference within the hypersphere identity metric space. Different from methods that performed recognition based on the recovered HR images [46, 57], our method is more like [18] that firstly learns the features, but differs in three aspects: 1) Our *non-linear and compact* features learned from the deep CNNs are more powerful than the linear features in [46, 57]. 2) Our FAN focuses on the disentangled identity features and thus can fully leverage both paired and unpaired data. 3) Our FAN well handles LR *surveillance* faces.

Face Normalization It is widely assumed in prior work [3, 6, 39] that the appearance of a face is influenced by two factors: identity and intra-class (or *non-identity*) variation. Normally face normalization is a general task of generating an identity-preserved face while removing other non-identity variation including pose, expression, illumination, and resolution. Most prior works of face normalization have focused on specifically removing pose variation, i.e., face frontalization [21, 42, 53], expression variation [1, 63], or illumination variation [7]. Other works [35, 63] perform pose and expression normalization. In contrast, our work mainly focuses on resolution normalization by enhancing facial details, which also handles PIE implicitly. Motivated by the disentanglement-based face recognition approaches [16, 42], we incorporate disentangled feature learning for LR face normalization.

Our work is close to face SR. Early face SR work [2, 26, 48] adopt different types of machine learning algorithms. For example, Baker and Kanade [2] learned a prior on the spatial distribution of the image gradient for frontal faces. Yang et al. [48] assumed that facial landmarks can be accurately estimated from the LR face image, and incorporated the facial priors by using the mapping between specific facial components. Recently, deep CNN has shown its superiority for face SR. Zhou et al. [60] proposed a bi-channel CNN for faces with large appearance variations. Zhu et al. [12] super-resolved unaligned low-resolution faces in a task-alternating cascaded framework. More recently, several works [8, 55] adopt Generative Adversarial Networks (GAN) to generate photo-realistic face images. However, the above methods ignore the identity information during training, which is essential for human perception [57] and downstream FR tasks. In contrast, our method jointly learns

surveillance face recognition and normalization, which is based on the disentangled identity information. We compare with the most relevant FSR papers in Tab. 1.

3. Feature Adaptation Network

In this section, we first give an overview of FAN (Section 3.1) that consists of two stages. A feature disentanglement framework is introduced in Section 3.2, which aims to disentangle identity features with other factors of variations. Then we propose the feature adaptation with random scale augmentation in Section 3.3 to learn resolution-robust identity features from both paired and unpaired training data.

3.1. Framework Overview

The goals of our work are two-folds: 1) resolution-robust face recognition; 2) identity-preserved and resolution-enhanced face normalization. We propose to learn a disentangled representation to achieve both tasks. Performing face normalization from disentangled features enables identity supervision on both the disentangled features and the normalized faces, in contrast to previous SR work where identity supervision is only applied to the super-resolved faces. Such identity supervisions allow us to leverage real *unpaired* HR and LR faces without pixel correspondence. This is crucial for tackling surveillance scenario where paired images are usually unavailable in a large scale.

As shown in Fig. 3, our method consists of two stages: disentangled feature learning (green components) and feature adaptation (orange components). Feature disentanglement has been successfully applied to face recognition and face synthesis [31, 42]. A disentangled representation is not only generative for face synthesis but also discriminative for face recognition. In this stage, we train our feature disentanglement framework with HR faces. A face image is encoded into identity and non-identity features, which are combined to generate the input image.

In the second stage, we fix all models in the disentanglement framework and perform feature adaptation with HR-LR input images that can be either *paired* or *unpaired*. A LR feature encoder is learned to extract discriminative identity features from LR faces. The disentanglement framework provides strong supervisions in the feature adaptation process. To achieve resolution-robust recognition, we propose Random Scale Augmentation (RSA) to overcome the drawbacks of fixed-scale SR in previous work.

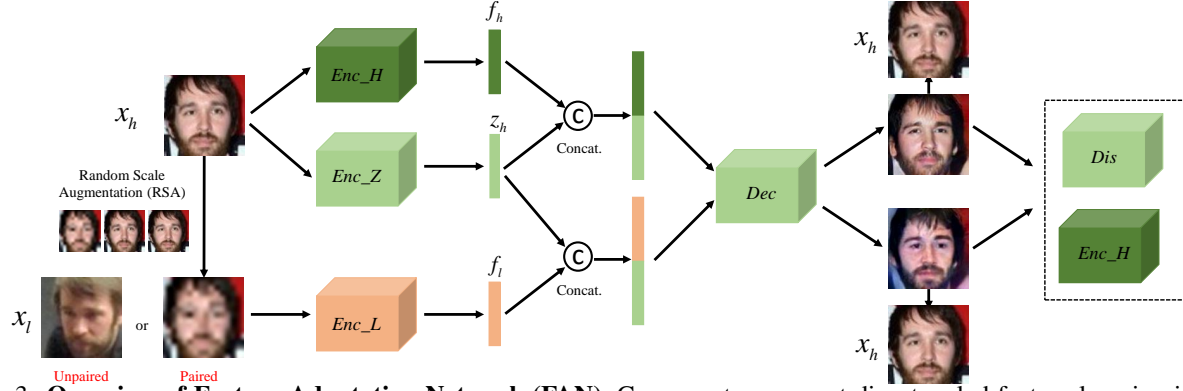


Figure 3: **Overview of Feature Adaptation Network (FAN).** Green parts represent disentangled feature learning in stage one, where dark green and light green are the two specific steps. First, Enc_H (dark green) is a HR identity encoder that is pre-trained and fixed. Second, Enc_Z , Dec , and Dis are trained for feature disentanglement. Orange parts represent the feature adaptation in stage two where a LR identity encoder is learned with all other models (green) fixed.

3.2. Disentangled Feature Learning

Our feature disentanglement framework consists of five modules: an identity feature encoder Enc_H , a non-identity feature encoder Enc_Z , a decoder Dec , a linear classifier FC (omitted from Fig. 3 for clarity), and a discriminator Dis . To disentangle identity features from the non-identity variations, we perform a two-step feature disentanglement motivated by [16], but differs in three aspects.

In the first step, a state-of-the-art face recognition model is trained with HR and down-sampled LR images using standard softmax loss and $m\text{-}L_2$ regularization [54]. We denote the trained feature encoder as Enc_H , which remains fixed for all later stages to provide encoded identity features $\mathbf{f}_h = Enc_H(\mathbf{x}_h)$ from a HR input \mathbf{x}_h . In the second step, we aim to learn non-identity features $\mathbf{z}_h = Enc_Z(\mathbf{x}_h)$ by performing adversarial training and image reconstruction.

The first difference to [16] is the loss for \mathbf{z}_h . [16] minimizes the identity classification loss, which we found to be unstable during training as it is unbounded. Instead, we propose to enforce the non-identity features to be evenly classified to all identities to make \mathbf{z}_h identity *unrelated* [31].

$$L_z = \|FC(\mathbf{z}_h) - \mathbf{y}_z\|_2^2, \quad (1)$$

where $\mathbf{y}_z = [\frac{1}{N_D}, \dots, \frac{1}{N_D}] \in \mathbb{R}^{N_D}$ and N_D is the total number of identities in the training set. The gradient of this loss is used to update only Enc_Z but not FC .

The disentangled features are combined to generate a face image $\mathbf{x}'_h = Dec(\mathbf{f}_h, \mathbf{z}_h)$ with the goal of recovering the input: $L_{dec} = \|\mathbf{x}'_h - \mathbf{x}_h\|_2^2$. As \mathbf{f}_h is discriminative for face recognition, the non-identity components will be discarded from \mathbf{f}_h in the first step. The reconstruction will encourage Enc_Z to encode non-identity features \mathbf{z}_h that is complimentary to \mathbf{f}_h in order to recover the input face.

The second difference to [16] is that we employ an identity similarity regularization and GAN-based discriminator loss to impose identity similarity and improve visual quality




Figure 4: **Our feature disentanglement learning performs two kinds of reconstructions.** Top row denotes the reconstruction from the identity and non-identity features where we supervise on both pixel domain reconstruction (L_{dec}) and feature level regularization (L_{id} and L_{gan}). Bottom row denotes the reconstruction from the identity features where only feature level regularizations are used.

of the generated faces. Specifically, for identity loss we use Enc_H to extract features and regularize the feature distance: $L_{id} = \|\mathbf{f}_h - \mathbf{f}_{h'}\|_2^2$. For GAN-based discriminator loss, we use standard binary cross entropy classification loss that is omitted here for clarity.

The third difference to [16] is that we perform an additional reconstruction to generate a face image from the identity features alone: $\mathbf{x}'_{h0} = Dec(\mathbf{f}_h, \mathbf{0})$ where $\mathbf{0}$ represents a vector of 0 that is of the same dimension as \mathbf{z}_h . As the non-identity part is given as $\mathbf{0}$, we expect the generated face to be an identity-preserved and *normalized* face without variations in non-identity factors such as PIE. As there is no ground-truth target image, we impose identity and GAN loss respectively. This process is illustrated in Fig. 4.

This additional reconstruction has two benefits. First, it encourages the identity features alone to synthesize an identity-preserved face, which in turn prevents the non-identity features to encode identity information. Therefore, it results in better disentanglement. Second, the ability of Dec to reconstruct an identity-preserved face from the iden-

tity features alone is useful in the testing stage for enhancing facial details, as will be shown in Section 3.3.

3.3. Paired and Unpaired Feature Adaptation

In the first stage, we have learned disentangled representation and image reconstruction from HR images. The reason to train *Dec* with HR only is to force *Dec* to generate HR images, which is the goal to enhance resolution for face normalization. However, this framework will not work well for LR inputs. Therefore, we propose a feature adaptation scheme to learn a LR Encoder *Enc_L* for LR face recognition and normalization. We aim to learn a feature extractor that works well for input faces with various resolutions.

Training with paired data In conventional FSR, it is common to down-sample the HR faces to LR versions with a few fixed scales and use the HR images to supervise FSR. Such methods cannot well handle the various resolutions in real-world surveillance imagery. To solve this issue, we propose Random Scale Augmentation (RSA). Given a HR input $\mathbf{x}_h \in \mathbb{R}^{N_h \times N_h}$, we down-sample the image to a random resolution to obtain $\mathbf{x}_l \in \mathbb{R}^{K \times K}$, where $K \in [N_l, N_h]$ and N_l is the lowest pixel resolution we care about (e.g. 8). We call the HR images and down-sampled LR counterparts as paired data as they have pixel-to-pixel correspondence.

Training with unpaired data Unpaired data represents HR and LR face images from the same subject but do not have pixel-wise correspondence. As shown in Fig. 2, the distribution of paired data is very far away from that of unpaired data. However, conventional FSR cannot take advantage of such unpaired data. Fortunately, FAN can well handle such case as we conduct face normalization from the disentangled features. As shown in Fig. 3, FAN is suitable for both paired and unpaired training because it adapts the feature distributions between LR and HR images. We perform both feature-level and image-level similarity supervisions.

Specifically, given a LR face image \mathbf{x}_l that is obtained from either random down-sampling or unpaired source, we resize \mathbf{x}_l to the same dimension as \mathbf{x}_h using bicubic interpolation. Then we use *Enc_L* to extract identity features \mathbf{f}_l , which is regularized to be similar to the disentangled features of the corresponding HR input image:

$$L_{enc} = \|\text{Enc}_L(\mathbf{x}_l) - \text{Enc}_H(\mathbf{x}_h)\|_2^2. \quad (2)$$

This feature-level regularization adapts the features of LR images to the HR images in the disentangled feature space. The second regularization is defined in the recovered face image space. Recall that in the first stage *Dec* is trained to generate a HR image from the identity and non-identity features of a HR input face. If *Enc_L* can encode identity-preserved features, such features can replace the original HR identity features \mathbf{f}_h to recover the HR input face. Thus, we impose an image-level regularization:

$$L_{enc,dec} = \|\text{Dec}(\mathbf{f}_l, \text{Enc}_Z(\mathbf{x}_h)) - \mathbf{x}_h\|_2^2. \quad (3)$$

As the non-identity features encoded from the HR image contributes to generating the output, the original HR can

Table 2: **Detailed training steps of FAN.** The first stage (1.*) involves a two-step disentanglement learning process. The second stage (1.2) is the feature adaptation process.

Stage	Input	Training Models	Fixed Models
1.1	HR+LR	<i>Enc_H</i>	—
1.2	HR	<i>Enc_Z, FC, Dec, Dis</i>	<i>Enc_H</i>
2	HR+LR	<i>Enc_L</i>	<i>Enc_H, Enc_Z, Dec, Dis</i>

be used as the target for supervision. This formulation is fundamentally different to all previous face SR methods that cannot impose pixelwise supervision from unpaired data.

Both feature-level and image-level regularizations will enforce *Enc_L* to learn robust identity features \mathbf{f}_l . By varying the resolutions of the inputs, \mathbf{f}_l is resolution robust. Moreover, we also encourage the generated output to be realistic and identity preserving using the pre-trained discriminator *Dis* and *Enc_H*. The detailed training steps are summarized in Tab. 2. First, we train *Enc_H* using HR and down-sampled LR images in stage 1.1. Second, we train the second step feature disentanglement using HR images only in stage 1.2 by fixing *Enc_H*. Third, we train *Enc_L* with all other models fixed in stage 2.

Inference In the testing stage, we extract identity features $\mathbf{f}_l = \text{Enc}_L(\mathbf{x}_l)$ from a LR input for face recognition, and can further perform face normalization via *Dec*($\mathbf{f}_l, \mathbf{0}$) by setting the non-identity component to $\mathbf{0}$. Thus we do not require HR images during inference.

4. Experiments

4.1. Implementation Details

Datasets We conduct extensive experiments on several datasets including refined MS-Celeb-1M (MS1M) [11] for training, and LFW [19], SCface [13], QMUL-SurvFace [9] and WIDER FACE [50] for testing. Refined MS1M, a cleaned version of the original dataset [15], contains 3.8M images of 85K identities. For LFW, we down-sample the 6,000 face pairs to low resolution and adopt the standard evaluation protocol. SCface consists of HR (gallery) and LR (probe) face images of 130 subjects. Following [33], 50 subjects are used for fine-tuning and 80 subjects are for testing. We conduct face identification where HR images are used as the gallery set and LR images with different resolutions (captured at three distances: 4.2m for d_1 , 2.6m for d_2 and 1.0m for d_3) form the probe set. QMUL-SurvFace consists of very-low resolution face images captured under surveillance cameras, and is used for face verification.

Training Setting Following [58], we use five facial landmarks (eye centers, nose tip and mouth corners) to align a face image to 128×128 . We uniformly re-size the input LR images to a fixed size of 128×128 by bicubic interpolation, which makes our method suitable for the proposed RSA strategy. Our framework is implemented with the Torch7 toolbox [10]. Our *Enc_H* and *Enc_L* are based on



Figure 5: **Feature disentanglement visualization on LFW.** We show input face x_h (a), reconstructed face from identity features $Dec(f_h, 0)$ (c), non-identity features $Dec(0, z_h)$ (d), and both (b). Feature disentanglement effectively normalizes faces in (c), and preserves pose(1,2), illumination(3), expression(4), and occlusions(5,6) in (d).

ResNet-50 [17], Enc_Z , Dec , and Dis are similar to [42]. We train stages 1.1 and 1.2 with a learning rate of $2e^{-4}$ for 12 and 8 epochs respectively. Stage 2 is trained with a learning rate of $2e^{-5}$ for 6 epochs. We use Adam optimization [24]. For SCface experiments, we finetune Enc_L with refined-MS1M (paired) and SCface (unpaired) training set for 1,000 iterations with a learning rate of $1e^{-5}$.

4.2. Ablation Study

Effects of Disentangled Feature Learning First, we evaluate the effects of disentangled feature learning by visualizing the disentangled identity features f_h and non-identity features z_h through our trained Dec . As shown in Fig. 5, the fusing of f_h and z_h can successfully recover the original image. The identity features alone can generate an identity-preserved frontal face, while the PIE variations and background information are captured in the non-identity features. This suggests our framework effectively disentangles identity and non-identity features.

Our feature disentanglement framework can also be applied for feature transfer. Given two images from either the same or different subjects, our model generates identity features as f_{h_1} , f_{h_2} and non-identity features as z_{h_1} , z_{h_2} . We perform feature transfer from one image to the other as: $Dec(f_{h_1}, z_{h_2})$ and $Dec(f_{h_2}, z_{h_1})$. Fig. 6 shows some examples where our feature transfer can keep the original image’s identity and change the attributes (PIE) accordingly.

Effects of Joint Learning We conduct verification tests to demonstrate the effects of jointly learning both face recognition and normalization, compared with Enc_H that is only trained for face recognition, and state-of-the-art SR methods such as VDSR [23], SRResNet [28] and FSRNet/FSRGAN [8], that are only trained for face hallucination. With the standard verification protocol of LFW, we




Figure 6: **Feature transfer visualization between two images of same subject (a) or different subjects (b, c).** In each set of examples, the top row shows the original images and the bottom shows the transferred images. The transferred image keeps the source’s (the above image) identity and the target’s (the diagonal image) attributes.

Table 3: **Face verification on super-resolved / normalized faces from LFW.** The Enc_H achieves 99.5% on faces with the original resolutions.

Method	8×	RSA
	Acc./PSNR	Acc./PSNR
Enc_H	86.6%/-	70.8%/-
VDSR (CVPR’16) + Enc_H	85.7%/26.63	69.4%/25.49
SRResNet (CVPR’17) + Enc_H	86.2%/27.46	68.6%/24.85
FSRNet (CVPR’18) + Enc_H	89.7%/28.27	69.4%/25.25
FSRGAN (CVPR’18) + Enc_H	86.7%/26.36	67.0%/23.94
FAN (i.e., normalized image) + Enc_H	91.9%/-	76.8%/-
FAN (i.e., Enc_L)	95.2%/-	82.4%/-

down-sample the 6,000 test pairs to 16×16 with a $8 \times$ scale factor, and upscale back to 128×128 via bicubic interpolation. Unlike [57] that retrains recognition network on the super-resolved images, we directly use Enc_H to evaluate the accuracy (i.e., Acc.) of different face SR methods.

Since FAN handles face hallucination and frontalization simultaneously, it is not suitable to use pixel-wise evaluation metrics (e.g., PSNR) to evaluate our method. Instead, we compare face verification performance. From the results in Tab. 3, we have three observations. 1) The recognition performance drops significantly when processing LR faces. 2) After incorporating SR methods, FSRNet [8] achieves better verification (i.e., Acc.) and SR (i.e., PSNR) results than other SR methods. Our FAN achieves the best results among all SOTA SR methods. 3) It is more effective to learn identity features from LR input (Our Enc_L) than performing face hallucination and recognition on the normalized faces (Our normalized image + Enc_H).

Effects of Random Scale Augmentation We further conduct tests to demonstrate the effects of our RSA strategy under the same experimental setting as the above study. The only difference is that we randomly down-sample the 6,000 test pairs to the resolution interval between 8×8 and 32×32 , i.e., the scale factors from $16 \times$ to $4 \times$. The extreme LR 8×8 and various/unknown resolutions existed in test images are more common than a larger and fixed resolution case. As shown in Tab. 3, we have three observations. 1) Since im-



Figure 7: Face SR/normalization on WIDER FACE.



Figure 8: Face SR/normalization on heavily distorted images from WIDER FACE. Top: input images. Middle: Bulat *et al.* [4]. Bottom: our FAN (*i.e.*, $Dec(f_t, 0)$).

ages with much lower resolutions (*e.g.*, 8×8) are evaluated, the baseline Enc_H is much lower than the $8 \times$ case. 2) Since all of the SR methods are trained with specific scales (*e.g.*, $8 \times$), they are not suitable for test images with varying resolutions. 3) Our FAN is much better than other methods in this case, which demonstrates the effects of RSA and our FAN is more practical in real-world scenarios.

4.3. Comparisons with SOTA Methods

In this section, we conduct extensive comparisons with SOTA methods both quantitatively and qualitatively. First, we show the hallucination ability of our method on recovering HR faces from heavily distorted faces on WIDER FACE [50]. Second, we demonstrate the ability for preserving the identity information via verification/identification comparisons on QMUL-SurFace [9] and SCface [13].

Comparisons on WIDER FACE First, we qualitatively compare our FAN with CycleGAN [61], SRGAN [28], Wavelet-SRNet [20], FSRNet [8] and Bulat *et al.* [4] on WIDER FACE. Figure 7 illustrates the recovered images, where the results of the competitors are all imported directly from [4]. As we can see, both our FAN and [4] perform well on recovering valid faces in these cases. However, there are two differences between ours and the competitors. First, FAN is trained on refined-MS1M for joint face recognition and normalization, and we do **NOT** finetune our model with the WIDER FACE data, while the competitors are directly



Figure 9: Face SR/normalization on the verification set of QMUL-SurFace. Top 2 rows: a face pair of the same subject. Bottom 2 rows: a face pair of different subjects.

trained with $\sim 50K$ images from WIDER FACE. Second, FAN generates 128×128 HR images, while the competitors super-resolve to 64×64 , a relatively **easier** task.

We further test our method on some heavily distorted images that [4] fails to recover meaningful HR faces, and show the results in Fig. 8. Thank to our powerful encoder and decoder networks, even when dealing with extremely low-quality faces, our method still recovers valid faces that are much clearer than [4]. It is an open question that whether we shall recover a face from a heavily distorted image. To address this issue, one could estimate the face quality and determine when normalization should be applied, which is not in the scope of our work or most prior face SR methods but a good direction for future work.

Comparisons on QMUL-SurFace QMUL-SurFace includes very LR faces captured with surveillance cameras. It is a very challenging dataset as most of the faces are hardly visible. We compare our framework with SOTA face SR methods [8, 23, 28] to evaluate the performance on real world surveillance data. As shown in Fig. 9, previous works struggle to recover the high-frequency information from the input LR faces. In contrast, our FAN can consistently generate a high-quality frontal face that recovers identity information. In addition to this qualitative evaluation, we also conduct face verification evaluation on the super-resolved/normalized face images. As shown in Tab. 4, our FAN performs better than previous SR methods. We have also evaluated the performance of using Enc_L to extract features directly from the LR inputs. The results in Tab. 4 indicates that it is more effective to learn features rather than performing super-resolution for LR face recognition, which is consistent with the observation in Tab. 3.

Comparisons on SCface SCface defines face identification with unpaired HR and LR faces. It mimics the real-world surveillance watch-list problem, where the gallery contains HR faces and the probe consists of LR faces captured from surveillance cameras. The HR and LR images do not have

Table 4: **Face verification results on QMUL-Surv super-resolved faces evaluated with Enc_H except the last row.**

Method	TAR(%)@FAR				AUC	Mean Acc (%)
	30%	10%	1%	0.1%		
VDSR [23]	61.03	35.32	8.89	3.10	71.02	65.64
SRResNet [28]	61.81	34.03	8.36	2.07	71.00	65.94
FSRNet [8]	59.92	33.10	7.84	1.93	70.09	64.96
FSRGAN [8]	56.03	30.91	8.45	2.66	67.93	63.06
FAN (norm. face)	62.31	36.64	11.89	3.70	71.66	66.32
FAN (Enc_L)	71.30	44.59	12.94	2.75	76.94	70.88

Table 5: **Rank-1 performance of face identification on SCface testing set.** ‘-FT’ means fine-tuning with SCface training set. Most compared results are cited from [33] except ArcFace that we pretrained on refined MS1M.

Distance →	d1	d2	d3	avg.
LDMDs [49]	62.7	70.7	65.5	66.3
RICNN [56]	23.0	66.0	74.0	54.3
LightCNN [47]	35.8	79.0	93.8	69.5
Center Loss [44]	36.3	81.8	94.3	70.8
ArcFace (ResNet50) [11]	48.0	92.0	99.3	79.8
LightCNN-FT	49.0	83.8	93.5	75.4
Center Loss-FT	54.8	86.3	95.8	79.0
ArcFace (ResNet50)-FT	67.3	93.5	98.0	86.3
DCR-FT [33]	73.3	93.5	98.0	88.3
FAN	62.0	90.0	94.8	82.3
FAN-FT (no RSA)	68.5	92.3	97.8	86.2
FAN-FT (no Dec)	73.0	94.0	97.8	88.3
FAN-FT	77.5	95.0	98.3	90.3

correspondence in the pixel domain, which is difficult for previous face SR methods that requires pixel-level correspondence. Fortunately, it is not a problem for our FAN as we regularize the model training in the disentangled feature space. Following [33], we conduct experiments on the day-time data only. The first 80 subjects are used for testing and the rest 50 subjects are for fine-tuning the second stage of our method. In addition to the unpaired HR and LR images in the training set, we also perform RSA to generate LR images from the HR images for model fine-tuning.

We mainly compare with DCR [33] as it achieved SOTA results on SCface. As far as we know, almost all SOTA face recognition methods have not evaluated on SCface. For fair comparison, we implemented ArcFace [11] using the same backbone and also finetuned on SCface training set. As shown in Tab. 5, our FAN achieves the best results among all other methods that are not finetuned on SCface. After finetuning on SCface with RSA, we achieve new SOTA results. Note that DCR proposed to use decoupled training that learns feature mapping for faces at *each* resolution, and the resolution information is assumed to be given in the testing stage. However, such resolution information is often unavailable in practice. Nevertheless, our method still outperforms DCR by a large margin even though we do **NOT** use the resolution information at the testing stage. From the




Figure 10: **Face normalization on SCface testing set.** (a) input images at three resolutions (d_1 , d_2 , d_3). (b) normalized faces generated by Enc_H and Dec . (c) normalized faces generated by Enc_L and Dec . (d) HR gallery images. The number on top-left corner indicates the feature distance between the input / normalized face and the gallery HR face.

gap between FAN-FT and FAN-FT (no RSA), we can see the effectiveness of RSA for surveillance FR. We also conducted ablative experiments by removing the Dec . No Dec means that we only perform feature-level similarity regularization in the second stage. The results of FAN-FT (no Dec) suggests that joint learning of face normalization can help feature learning for FR.

Figure 10 shows our face normalization results on SCface testing set. Our method can generate high-quality face images that recover the identity information from the input faces with various resolutions. The comparison between the results generated by Enc_H and Enc_L validates the effectiveness of feature adaptation in our second stage both quantitatively and qualitatively.

Time Complexity Our framework adopts an encoder-decoder (*i.e.*, Enc_L and Dec) structure, it takes ~ 0.011 s to extract *compact* identity features, and another ~ 0.005 s to recover a 128×128 HR image on Titan X GPU, which is comparable to 0.012s for FSRNet [8] on the same hardware, and much faster than 3.84s for CBN [62], 8 min. for [30] and 15 min. for [22]. In general, compared with the SOTA methods, our FAN is a better choice in surveillance scenarios considering both visualization and recognition.

5. Conclusions

This paper proposes a Feature Adaptation Network (FAN) for surveillance face recognition and normalization. Despite the great improvement in face recognition and super-resolution, the applications in surveillance scenario is less studied. We aim to bridge this gap. FAN consists of two stages: feature disentanglement learning and feature adaptation. By first disentangling the face features into identity and non-identity components, it enables our adaptation network to impose both feature-level and image-level sim-

ilarity regularizations. Such framework is suitable for both paired and unpaired training, which overcomes the limit by previous face SR methods that require paired training data. The proposed Random Scale Augmentation (RSA) is very effective in handling the various resolutions in surveillance imagery. We achieved SOTA face recognition and normalization results even from very low quality inputs.

References

- [1] B. Amberg, R. Knothe, and T. Vetter. Expression invariant 3d face recognition with a morphable model. In *FG*, 2008. 3
- [2] S. Baker and T. Kanade. Hallucinating faces. In *FG*, 2000. 3
- [3] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(3):711–720, 1997. 3
- [4] A. Bulat, J. Yang, and G. Tzimiropoulos. To learn image super-resolution, use a gan to learn how to do image degradation first. In *ECCV*, 2018. 1, 2, 3, 7
- [5] J. Cai, H. Zeng, H. Yong, Z. Cao, and L. Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *ICCV*, 2019. 2
- [6] D. Chen, X. Cao, L. Wang, F. Wen, and J. Sun. Bayesian face revisited: A joint formulation. In *ECCV*, 2012. 3
- [7] W. Chen, M. J. Er, and S. Wu. Illumination compensation and normalization for robust face recognition using discrete cosine transform in logarithm domain. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 36(2):458–466, 2006. 3
- [8] Y. Chen, Y. Tai, X. Liu, C. Shen, and J. Yang. FSRNet: End-to-End learning face super-resolution with facial priors. In *CVPR*, 2018. 1, 2, 3, 6, 7, 8
- [9] Z. Cheng, X. Zhu, and S. Gong. Surveillance face recognition challenge. *arXiv preprint arXiv:1804.09691*, 2018. 2, 3, 5, 7
- [10] R. Collobert, K. Kavukcuoglu, and C. Farabet. Torch7: A matlab-like environment for machine learning. In *NIPS Workshop*, 2011. 5
- [11] J. Deng, J. Guo, N. Xue, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 1, 2, 5, 8
- [12] C. Dong, C. Loy, K. He, and X. Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 38(2):295–307, 2016. 3
- [13] M. Grgic, K. Delac, and S. Grgic. SCface - surveillance cameras face database. *Multimedia Tools and Applications*, 51(3):863–879, 2011. 2, 3, 5, 7
- [14] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Guide to the cmu multi-pie database. Technical report, Carnegie Mellon University. 3
- [15] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao. Ms-Celeb-1M: A dataset and benchmark for large-scale face recognition. In *ECCV*, 2016. 2, 5
- [16] N. Hadad, L. Wolf, and M. Shahar. A two-step disentanglement method. In *CVPR*, 2018. 3, 4
- [17] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [18] P. H. Hennings-Yeomans, S. Baker, and B. V. Kumar. Simultaneous super-resolution and feature extraction for recognition of low-resolution faces. In *CVPR*, 2008. 2, 3
- [19] 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. 1, 2, 3, 5
- [20] H. Huang, R. He, Z. Sun, and T. Tan. Wavelet-SRNet: A wavelet-based CNN for multi-scale face super resolution. In *ICCV*, 2017. 7
- [21] R. Huang, S. Zhang, T. Li, and R. He. Beyond face rotation: Global and local perception gan for photorealistic and identity preserving frontal view synthesis. In *ICCV*, 2017. 3
- [22] Y. Jin and C. Bouganis. Robust multi-image based blind face hallucination. In *CVPR*, 2015. 8
- [23] J. Kim, J. K. Lee, and K. M. Lee. Accurate image super-resolution using very deep convolutional networks. In *CVPR*, 2016. 6, 7, 8
- [24] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [25] B. F. Klare, B. Klein, E. Taborsky, A. Blanton, J. Cheney, K. Allen, P. Grother, A. Mah, and A. K. Jain. Pushing the frontiers of unconstrained face detection and recognition: larpa janus benchmark a. In *CVPR*, 2015. 1, 2
- [26] S. Kolouri and G. K. Rohde. Transport-Based single frame super resolution of very low resolution face images. In *CVPR*, 2015. 3
- [27] V. Le, J. Brandt, Z. Lin, L. Boudev, and T. S. Huang. Interactive facial feature localization. In *ECCV*, 2012. 3
- [28] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, 2017. 6, 7, 8
- [29] P. Li, J. Brogan, and P. J. Flynn. Toward facial re-identification: Experiments with data from an operational surveillance camera plant. In *BTAS*, 2016. 2
- [30] C. Liu, H.-Y. Shum, and W. T. Freeman. Face hallucination: Theory and practice. *International Journal of Computer Vision (IJCV)*, 75(1):115–134, 2007. 8
- [31] Y. Liu, F. Wei, J. Shao, L. Sheng, J. Yan, and X. Wang. Exploring disentangled feature representation beyond face identification. In *CVPR*, 2018. 3, 4
- [32] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. 3
- [33] Z. Lu, X. Jiang, and A. Kot. Deep coupled resnet for low-resolution face recognition. *IEEE Signal Processing Letters*, 25(4):526–530, 2018. 5, 8
- [34] P. J. Phillips, P. J. Flynn, T. Scruggs, K. Bowyer, J. Chang, K. Hoffman, J. Marques, J. Min, and W. Worek. Overview of the face recognition grand challenge. In *CVPR*, 2005. 3
- [35] Y. Qian, W. Deng, and J. Hu. Unsupervised face normalization with extreme pose and expression in the wild. In *CVPR*, 2019. 3

- [36] M. S. Rad, B. Bozorgtabar, U.-V. Marti, M. Basler, H. K. Ekenel, and J.-P. Thiran. Srobb: Targeted perceptual loss for single image super-resolution. In *ICCV*, 2019. 2
- [37] S. Sengupta, J.-C. Chen, C. Castillo, V. M. Patel, R. Chellappa, and D. Jacobs. Frontal to profile face verification in the wild. In *WACV*, 2016. 2
- [38] S. Shekhar, V. M. Patel, and R. Chellappa. Synthesis-based recognition of low resolution faces. In *IJCB*, 2011. 2
- [39] J. Susskind, G. Hinton, R. Memisevic, and M. Pollefeys. Modeling the joint density of two images under a variety of transformations. In *CVPR*, 2011. 3
- [40] Y. Tai, J. Yang, and X. Liu. Image super-resolution via deep recursive residual network. In *CVPR*, 2017. 2
- [41] Y. Tai, J. Yang, X. Liu, and C. Xu. MemNet: A persistent memory network for image restoration. In *ICCV*, 2017. 2
- [42] L. Tran, X. Yin, and X. Liu. Disentangled representation learning GAN for pose-invariant face recognition. In *CVPR*, 2017. 2, 3, 6
- [43] H. Wang, Y. Wang, Z. Zhou, X. Ji, D. Gong, J. Zhou, Z. Li, and W. Liu. Cosface: Large margin cosine loss for deep face recognition. In *CVPR*, 2018. 1, 2
- [44] Y. Wen, K. Zhang, Z. Li, and Y. Qiao. A discriminative feature learning approach for deep face recognition. In *ECCV*, 2016. 1, 2, 8
- [45] L. Wolf, T. Hassner, and I. Maoz. Face recognition in unconstrained videos with matched background similarity. In *CVPR*, 2011. 3
- [46] J. Wu, S. Ding, W. Xu, and H. Chao. Deep joint face hallucination and recognition. *arXiv:1611.08091v1*, 2016. 2, 3
- [47] X. Wu, R. He, Z. Sun, and T. Tan. A light CNN for deep face representation with noisy labels. *IEEE Transactions on Information Forensics and Security*, 13(11):2884–2896, 2018. 8
- [48] C.-Y. Yang, S. Liu, and M.-H. Yang. Structured face hallucination. In *CVPR*, 2013. 3
- [49] 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. 8
- [50] S. Yang, P. Luo, C. Loy, and X. Tang. Wider face: A face detection benchmark. In *CVPR*, 2016. 3, 5, 7
- [51] B. Yin, L. Tran, H. Li, X. Shen, and X. Liu. Towards interpretable face recognition. In *ICCV*, 2019. 2
- [52] X. Yin and X. Liu. Multi-task convolutional neural network for pose-invariant face recognition. *IEEE Transactions on Image Processing (TIP)*, 27(2):964–975, 2018. 2
- [53] X. Yin, X. Yu, K. Sohn, X. Liu, and M. Chandraker. Towards large-pose face frontalization in the wild. In *ICCV*, 2017. 3
- [54] X. Yin, X. Yu, K. Sohn, X. Liu, and M. Chandraker. Feature transfer learning for face recognition with under-represented data. In *CVPR*, 2019. 4
- [55] X. Yu and F. Porikli. Ultra-resolving face images by discriminative generative networks. In *ECCV*, 2016. 2, 3
- [56] D. Zeng, H. Chen, and Q. Zhao. Towards resolution invariant face recognition in uncontrolled scenarios. In *ICB*, 2016. 8
- [57] K. Zhang, Z. Zhang, C.-W. Cheng, W. H. Hsu, Y. Qiao, W. Liu, and T. Zhang. Super-Identity convolutional neural network for face hallucination. In *ECCV*, 2018. 2, 3, 6
- [58] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, 23(10):1499–1503, 2016. 5
- [59] J. Zhao, L. Xiong, Y. Cheng, Y. Cheng, J. Li, L. Zhou, Y. Xu, J. Karlekar, S. Pranata, S. Shen, and et al. 3D-aided deep poseinvariant face recognition. In *IJCAI*, 2018. 2
- [60] E. Zhou, H. Fan, Z. Cao, Y. Jiang, and Q. Yin. Learning face hallucination in the wild. In *AAAI*, 2015. 3
- [61] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 7
- [62] S. Zhu, S. Liu, C. Loy, and X. Tang. Deep cascaded bi-network for face hallucination. In *ECCV*, 2016. 2, 8
- [63] X. Zhu, Z. Lei, J. Yan, D. Yi, and S. Z. Li. High-fidelity pose and expression normalization for face recognition in the wild. In *CVPR*, 2015. 2, 3