

Joint Super-Resolution and Alignment of Tiny Faces

Yu Yin,¹ Joseph P. Robinson,¹ Yulun Zhang,¹ Yun Fu^{1,2}

¹Department of Electrical and Computer Engineering, Northeastern University, Boston, MA

²Khoury College of Computer Science, Northeastern University, Boston, MA

{yin.yu1, robinson.jo}@husky.neu.edu, yulun100@gmail.com, yunfu@ece.neu.edu

Abstract

Super-resolution (SR) and landmark localization of tiny faces are highly correlated tasks. On the one hand, landmark localization could obtain higher accuracy with faces of high-resolution (HR). On the other hand, face SR would benefit from prior knowledge of facial attributes such as landmarks. Thus, we propose a joint alignment and SR network to simultaneously detect facial landmarks and super-resolve tiny faces. More specifically, a shared deep encoder is applied to extract features for both tasks by leveraging complementary information. To exploit representative power of the hierarchical encoder, intermediate layers of a shared feature extraction module are fused to form efficient feature representations. The fused features are then fed to task-specific modules to detect landmarks and super-resolve face images in parallel. Extensive experiments demonstrate that the proposed model significantly outperforms the state-of-the-art in both landmark localization and SR of faces. We show a large improvement for landmark localization of tiny faces (*i.e.*, 16×16). Furthermore, the proposed framework yields comparable results for landmark localization on low-resolution (LR) faces (*i.e.*, 64×64) to existing methods on HR (*i.e.*, 256×256). As for SR, the proposed method recovers sharper edges and more details from LR face images than other state-of-the-art methods, which we demonstrate qualitatively and quantitatively.

Introduction

Automatic face understanding is critical for problems in human perception (*e.g.*, super-resolution (SR) (Yu and Porikli 2016), visual understanding (Güçlütürk et al. 2017), and style transfer (Liu, Breuel, and Kautz 2017)) and applied machine vision (*e.g.*, landmark localization (Robinson et al. 2019), identity recognition (Wu et al. 2016), and face detection (Zhang et al. 2016)). Modern-day models for face-based tasks tend to breakdown when applied to images of low-resolution (LR). In practice, face-based systems are frequently confronted with such scenarios (*e.g.*, LR cameras used for surveillance (Yu and Porikli 2017)). Recent studies revealed that a decrease in resolution (*i.e.*, $< 30 \times 30$) yields an increase in error for models used for facial landmark localization (Bulat et al. 2018). To address this problem, face

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Comparison with SuperFAN (Bulat et al. 2018) and FSRNet (Chen et al. 2018). The proposed recovers sharper edges and finer details HR space (a). Also, (b) shows estimated landmarks superimposed on SR face, where Red marks true landmarks and green marks predicted landmarks.

SR, also known as face hallucination, aims to generate high-resolution (HR) faces from LR imagery (Liu, Shum, and Freeman 2007). The recovered faces then provide more detailed information (*e.g.*, sharper edges, clearer shapes, and finer skin details), and are often used for improved analysis and perception. However, most existing methods (*e.g.*, Superfan (Bulat and Tzimiropoulos 2018)) rely heavily on the quality of recovered images. Since SR methods usually suffer from blurriness, using SR images for face-related tasks can hinder the final prediction or conclusion.

On the other hand, facial prior knowledge can be used to recover SR faces of higher quality (Baker and Kanade 2000; Liu, Shum, and Freeman 2007). In problems of single image super-resolution (SISR), face SR utilizes prior knowledge to improve the accuracy of the inferred images and, thus, to yield results of higher quality. For example, one can leverage low-level information (*i.e.*, smoothness in color), facial heatmaps, and face parsing maps to provide additional mid-level information (*i.e.*, face structure) to recover sharper edges and shapes (Chen et al. 2018). Also, high-level information can be extracted with identity labels and other face

attributes (*e.g.*, gender, age, and pose), and then leveraged to reduce the ambiguity of the hallucinated faces (Yu et al. 2018; Lee et al. 2018). Hence, additional face information is beneficial for SR, and especially for tiny faces (*e.g.*, 16×16).

Previous work in face SR either super-resolved LR images using prior information (*e.g.*, FSRNet (Chen et al. 2018)) or directly localized the landmarks on the super-resolved images (*e.g.*, SuperFAN (Bulat and Tzimiropoulos 2018)). Figure 2 compares these frameworks with the proposed method. Specifically, SuperFAN only uses SR to help localize the landmarks of tiny faces, but not vice-versa. Besides, our model does not process the recovered SR output that suffers from blurriness, as we dedicate an encoding module to maximize the amount of information captured from LR faces. As for FSRNet, landmarks are only used as facial prior knowledge to super-resolve faces, which suffers from the same problem of detecting landmarks on a coarse, recovered SR image. Furthermore, SuperFAN and FSRNet address the two tasks separately, leading to redundant feature maps. Since both face SR and landmark localization tasks could benefit from one another, we aim to extract the maximum amount of information from LR faces by addressing the two tasks simultaneously. Thus, we propose a multi-task framework that allows these tasks to benefit from one another, which improves the performance in both tasks (see Figure1).

The main contribution of this paper are as follows:

1. In this paper, we propose a network that does SR and landmark detection on tiny faces jointly– a network we dubbed JASRNet¹. To the best of our knowledge, we are the first to train a multi-task model that jointly learns landmark localization and SR. Specifically, and unlike existing two-step approaches, we leverage the complementary information of the two tasks. This allows for more accurate landmark predictions to be made in LR space and improved reconstruction from LR-to-HR.
2. Novel deep feature extraction and fusion modules are used to maximize the amount of information captured from the LR faces, which is done at intermediate layers of the encoder to exploit the deep hierarchical machinery.
3. We show large improvements for both SR and landmark localization for tiny faces (*i.e.*, 16×16). Besides, our JASRNet yields results for landmark localization on LR faces (*i.e.*, 64×64) that are comparable to existing methods evaluated on the corresponding HR faces (*i.e.*, 256×256). Furthermore, the proposed method recovers HR faces with sharper edges and shapes compared with state-of-the-art methods for SR.

Related Work

Face super-resolution

Typical SISR methods do not benefit from facial prior information and can be utilized to super-resolve images of arbitrary type. By introducing face-specific information, Yu (Yu and Porikli 2016; 2017) proposed a GAN-based model to recover HR images from tiny faces of size 16×16 . Chen (Chen

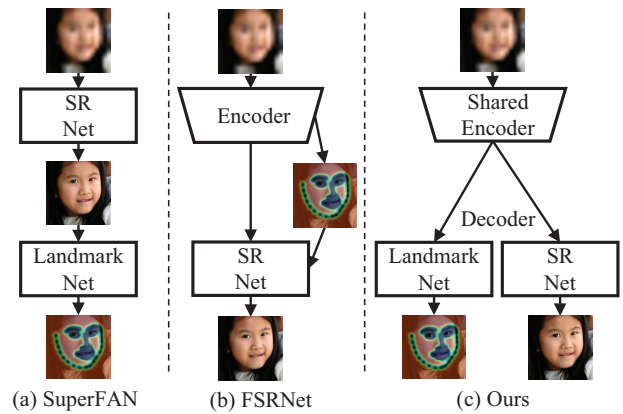


Figure 2: Graphical view. (a) SuperFAN (Bulat and Tzimiropoulos 2018) detects landmarks on super-resolved faces. (b) FSRNet (Chen et al. 2018) uses prior information for SR. (c) Our multi-task framework jointly learns landmark localization and SR, with tasks aiding one another.

et al. 2018) used a separate branch to estimate facial landmark heatmaps and parsing maps, which were then used as face-specific information to super-resolve tiny face images. FaceAttr (Yu et al. 2018) validated that knowledge of facial attributes can also significantly reduce the ambiguity in face SR. It is worth noting that our method not only utilizes facial prior information to super-resolve tiny face with better quality, but also achieves state-of-the-art performance on landmark alignment by benefiting from SR.

Face alignment

Modern-day approaches for face alignment have been successful on HR faces (Dong et al. 2018; Lv et al. 2017; Mo et al. 2019; Ranjan, Patel, and Chellappa 2019). However, most suffer from performance degradation with decreasing image resolution, especially with faces smaller than 30×30 (Bulat and Tzimiropoulos 2017). The first to address landmark detection on LR faces was SuperFAN (Bulat and Tzimiropoulos 2018), which super-resolved tiny faces, from which the output images were fed to a landmark localization model. Although the error of the landmark localization provides gradients to back-propagate through the SR module, it is, in essence, a 2-step process. We argue that the facial prior information is not fully utilized for SR. To address this problem, we present a novel synergistic multi-task framework that learns facial landmark localization and SR jointly.

Multi-task learning

Multi-task learning is commonly used to jointly address correlated tasks. HyperFace (Ranjan, Patel, and Chellappa 2019) proposed a multi-task learning framework for face detection, face alignment, gender recognition, and pose estimation. The joint learning tasks were based on regression or classification (*i.e.*, a special case of regression). Hence, similar architectures were adopted for all tasks. In our case, however, face SR and alignment are based on generation and regression, respectively. Thus, one of the main differences in

¹The code is available at: <https://github.com/YuYin1/JASRNet>.

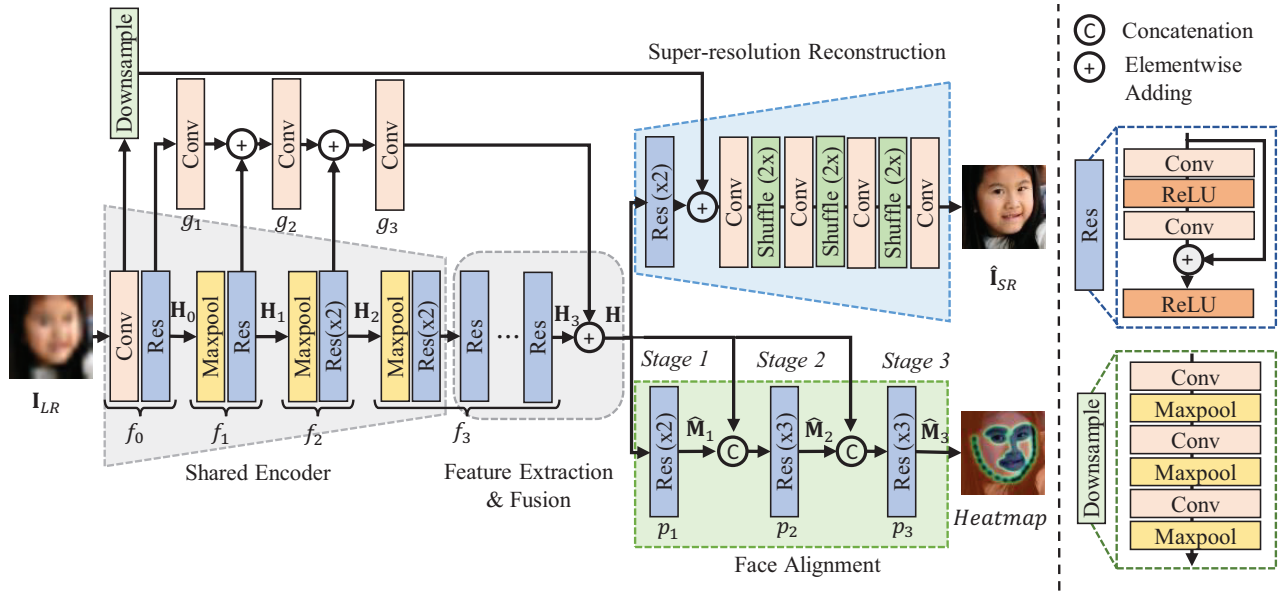


Figure 3: Architecture of the proposed JASNet. The shared encoder module is used for extracting shallow and shared features for both tasks. The deep feature extraction and fusion module is used for obtaining better feature representations. The other two modules are task-specific modules for super-resolution and face alignment, respectively.

architectures of the proposed from HyperFace is that we include specific modules for each task, while HyperFace used only fully connected layers after feature fusion.

Method

Super-resolution (SR) and landmark localization of tiny faces are highly correlated tasks. Both of them can benefit from each other. While previous work either uses SR to help align tiny faces or vice-versa, but not both. We argue that the amount of information extracted from LR image is not maximized when only one task is used to help the other. Hence, we propose a deep joint alignment and super-resolution network (JASNet) to model super-resolution and localize landmarks for tiny faces simultaneously, with information from both tasks boosting the performance of the other. As shown in Figure 3, the proposed JASNet consists of four parts: (1) a shared shallow encoder module is used for extracting shallow and shared features for both tasks; (2) a deep feature extraction and fusion, which is used for obtaining better feature representations; (3-4) task-specific modules for super-resolution and face alignment, respectively.

Let $\{\mathbf{I}_{LR}^{(i)}, \mathbf{I}_{HR}^{(i)}, \mathbf{M}^{(i)}\}_{i=1}^N$ be N training samples. The original LR faces $\mathbf{I}_{LR}^{(i)}$ are passed in to the shared encoder, which then feeds into the feature extraction module to extract features for both tasks. To exploit the representative power of different grains, the intermediate features of the shared encoder branch out to fuse with the output of deep feature extraction module. This feature fusion forms a more efficient feature representation, as later demonstrated as part of the ablation study. Carrying on, the fused features are fed to both task-specific modules. Thus, the super-resolved images $\hat{\mathbf{I}}_{SR}^{(i)}$ and the probability maps of the landmark estima-

tions $\hat{\mathbf{M}}^{(i)}$ are produced simultaneously.

Usually, there are sharper edges or sudden changes around the contour of facial component. For face alignment, the SR module recovers the image with better resolution, which, hence, helps the model detect more accurate landmarks. In parallel, the alignment module locates the edges and structure of the face, forcing more attention to the high-frequency content (*i.e.*, edges). Since both tasks, face SR and landmark localization, are suited to benefit from one another, the aim of this work is to exploit the amount of maximum information that can be extracted from the LR faces. This is done by combining the loss function of each task. For the SR task, the L_1 loss is minimized, as it can provide better convergence than L_2 (Lim et al. 2017; Zhang et al. 2018). For the alignment task, a L_2 heatmap loss is used, like in (Dong et al. 2018). Together, the loss function of JASNet can be expressed as

$$\begin{aligned}
 l &= l_{sr} + \alpha l_{heatmap} \\
 &= \frac{1}{N} \sum_{i=1}^N \{ \|\hat{\mathbf{I}}_{SR}^{(i)} - \mathbf{I}_{HR}^{(i)}\|_1 + \alpha \|\hat{\mathbf{M}}^{(i)} - \mathbf{M}^{(i)}\|_2 \}, \quad (1)
 \end{aligned}$$

where l denotes the total loss, l_{sr} and $l_{heatmap}$ denote the L_1 loss for super-resolution and the L_2 heatmap loss for alignment, respectively. The weight of $l_{heatmap}$ is α , and the estimated heatmap of the i^{th} image is $\hat{\mathbf{M}}^{(i)}$. As mentioned above, $\hat{\mathbf{I}}_{SR}^{(i)}$ is the super-resolved image recovered from $\mathbf{I}_{LR}^{(i)}$.

Shared feature extraction and fusion

Shallow encoder. Previous work in face SR and alignment usually addressed these two tasks separately, leading to redundant feature maps. To efficiently extract features from

LR images, a shared encoder is designed to extract shallow features that capture complementary information of the two tasks. It consists of a convolutional layer, a residual block (He et al. 2016), and then three transformations made-up of the maxpooling operation and residual blocks (Figure 3). Intermediate layers of the encoder are later fused for richer features in geometry and semantics.

All the convolution layers of JASRNet use kernels of size 3×3 , and each is followed by a ReLU layer. The number of channels are all set as 128, except for the last convolutional layers in both reconstruction and alignment module, which are set as 3 and the number of landmarks (namely 68 for 300W), respectively. There are three maxpooling layers in the network, each downsample the feature maps $\times 2$, which, in total, reduce the size of the feature maps by a factor of 8. The structure of the residual blocks is the same as in the original residual nets (ResNets) (He et al. 2016), except we omitted the batch normalization (BN) layers, as it reduces the variation of feature ranges: ResNets used for SISR (EDSR) performed best with all BN layers removed (Lim et al. 2017). Also, we found that BN layers slow down the speed to convergence of the network, while reducing its overall performance, which was especially true in the SR task. Since we aim to reserve the most information possible when passing through the shared encoder module (*i.e.*, during feature extraction), we follow EDSR (Lim et al. 2017) and remove all BN from the residual blocks.

Deep feature extraction and fusion. Deeper networks have shown to have a better performance in many computer vision tasks including SR (Bulat and Tzimiropoulos 2018; Chen et al. 2018; He et al. 2016; Lim et al. 2017; Tai et al. 2017). Increased depth was also a tactic used in this work. Shallow features extracted from the shared encoder are passed to the deep feature extraction module consisting of T residual blocks, with $T = 32$ in the reported experiments. A deeper network not only recovers sharper edges and shapes for super-resolved face images, but it also achieves a higher accuracy for landmark localization.

Inspired by Hyperface (Ranjan, Patel, and Chellappa 2019), we fused intermediate layers to exploit the representative power of features at different levels of the hierarchical model. Considering the similarity of features from adjacent layers, not all features of the shared encoder are fused to compose the new feature representation. Since each of the maxpooling layers downsample the feature map by a factor of 2, the output of the layers that precede each maxpooling layer branches out using skip connections, and are later fused to form richer features with geometry information. To match sizes of the feature maps, a 3×3 convolutional layer with stride 2 is applied to downsample fusing features by a factor of 2 for each maxpooling layer that is applied in parallel to the skip connection.

The outputs before the maxpooling are denoted as $\mathbf{H}_i (i \in \{0, 1, 2\})$; the output of the last residual block in feature extraction module is \mathbf{H}_3 (see Figure 3). Provided LR images \mathbf{I}_{LR} as input, we have

$$\begin{aligned} \mathbf{H}_0 &= f_0(\mathbf{I}_{LR}), \\ \mathbf{H}_i &= f_i(\mathbf{H}_{(i-1)}), i \in \{1, 2, 3\}, \end{aligned}$$

where $f_i(\cdot) (i \in \{0, 1, 2, 3\})$ transform the signal during feature extraction. Hence, f_0 is the mapping of the first convolution layers and residual blocks, f_1 and f_2 are the mappings of the first and second steps combining maxpooling and residual blocks, respectively, and f_3 is the mapping for the remaining residual blocks making up the feature extraction module.

Mathematically speaking, the fused features \mathbf{H} that is output can be founded as

$$\mathbf{H} = g_3(g_2(g_1(\mathbf{H}_0) + \mathbf{H}_1) + \mathbf{H}_2) + \mathbf{H}_3, \quad (2)$$

where the convolution operation $g_i(\cdot) (i \in \{1, 2, 3\})$ fuses intermediate features.

Task-specific modules

Super-resolution reconstruction. The *super-resolution reconstruction* module reconstructs the HR image from shared features of size 16×16 . First, shared feature maps are fed to two residual blocks to extract task-specific features. Next, 3 conv-layers, each of which are followed by pixel shuffle layers (Shi et al. 2016), upscale the feature maps $2 \times$ in size (*i.e.*, 16×16 to 128×128). Finally, a convolutional layer made-up of 3×3 filters to map from HR RGB image space.

Inspired by EDSR (Lim et al. 2017) and RDN (Zhang et al. 2018), the first and last residual blocks of the shared encoder and SR reconstruction module are linked by a large skip connection. This recovers HR images with finer details (*i.e.*, sharper edges and shapes). The skip connection directly provides low frequency information to the super-resolved images. Hence, it forces the network to focus on learning the high frequency information, opposed to low frequency information already provided. Since the output size of the first convolution layer is 128×128 , and the feature map size of the last residual block in reconstruction module is 16×16 , we downsample the 128×128 feature map $8 \times$ with 3 convolution and 3 maxpooling layers (see Figure 3).

Unlike SuperFAN, where the long skip connection is reported to have minimal impact on overall performance, our model largely benefits from the skip connection. This is because the features extracted includes high frequency information and, thus, is more efficient for recovering sharp and accurate edges. Furthermore, since super-resolution and face alignment share the deep features, a byproduct of this long skip connection also is boosted performance for the landmark localization task as well.

Face alignment. Like the *SR reconstruction* module, the shared features \mathbf{H} are fed through consecutive residual blocks to extract features specific to face alignment. Inspired by the successes of convolutional pose machines (CPM) (Wei et al. 2016) on face alignment, we also utilize the sequential framework made-up of residual blocks for estimating locations of landmarks. In the first stage, two residual blocks predict coarse heatmaps $\hat{\mathbf{M}}_1$. Then, in the second stage, the heatmaps $\hat{\mathbf{M}}_1$ predicted in the first stage are first concatenated with the feature maps \mathbf{H} , which are then fed to the second prediction module composed of three sequential residual blocks that predict heatmaps $\hat{\mathbf{M}}_2$. The third stage then concatenates feature maps \mathbf{H} and $\hat{\mathbf{M}}_2$ to produce final

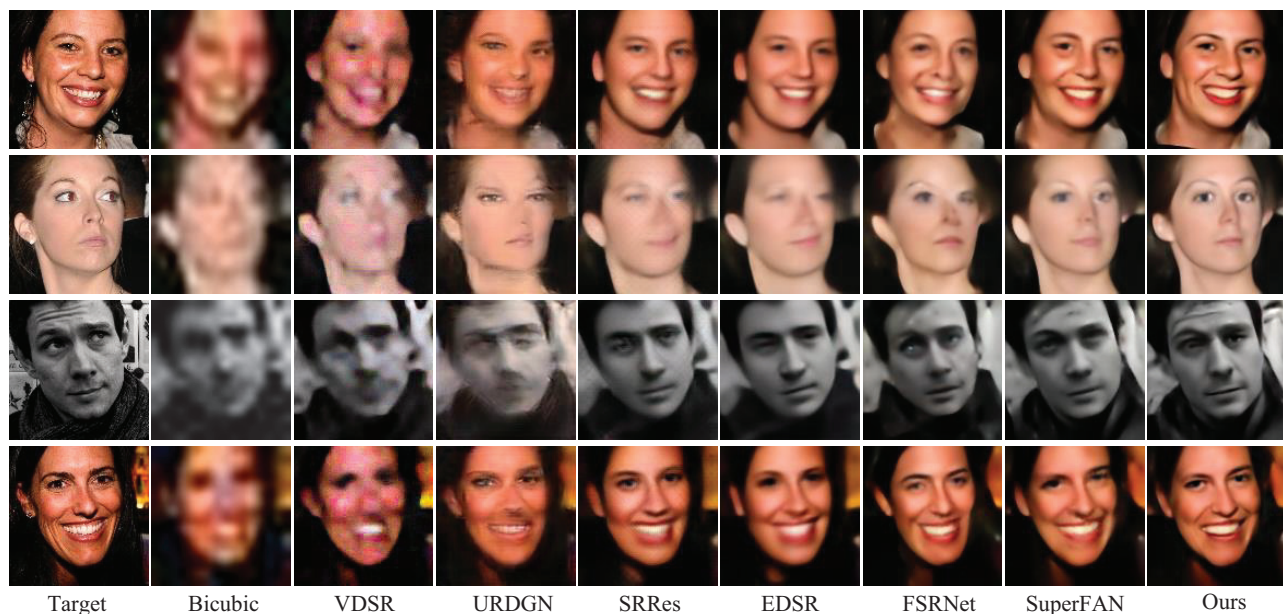


Figure 4: Visual results. Comparison of different super-resolution methods.

Table 1: Quantitative comparisons. PSNR/SSIM on 300W and HELEN.

	Bicubic	VDSR	URDGN	SRRes	EDSR	TDAE	FSRNet	SuperFAN	Ours
300W	21.36/0.594	21.80/0.558	21.97/0.617	23.30/0.669	23.47/0.658	21.12/0.547	23.05/0.678	23.13/0.691	23.69/0.711
HELEN	21.36/0.593	21.66/0.552	21.77/0.605	23.05/0.674	23.40/0.709	21.70/0.542	- / -	23.17/0.695	23.55/0.717

estimation \hat{M}_3 expressed as

$$\hat{M}_3 = p_3([\mathbf{H}, p_2([\mathbf{H}, p_1(\mathbf{H})]])], \quad (3)$$

where p_j maps the prediction modules, with $(j \in \{1, 2, 3\})$. Note that the size of the feature maps is constant throughout the face alignment module (*i.e.*, 16×16). During training, heatmap regression L_2 loss was used to localize landmarks, opposed to directly predicting pixel coordinates (x, y) . Thus, argmax is used to determine (x, y) from the predicted heatmaps in final stage (*i.e.*, \hat{M}_3). Specifically, the maximum value of each of the K heatmaps is found as the predicted landmarks (*i.e.*, $\text{argmax}_{(x,y)} \hat{M}_3$).

Experiments

We now review the experimental settings and results. Specifically, the datasets, implementation details, and metrics are first described. Then, we show results comparing with the state-of-the-art methods for the face SR and alignment task separately. Besides, we highlight the benefits of the proposed feature fusion and joint training. Finally, we conduct an ablation study as a deep-dive revealing the contributions of the components introduced in this work.

Experimental settings

Datasets. We evaluated the proposed approach on several datasets, which are listed as follows:

- **300W** (Sagonas et al. 2013; 2016) consists of 3,837 face images with 68 landmarks. We used the same training set as (Lv et al. 2017; Zhu et al. 2015). Subsets of 300W were evaluated: *common* and *challenge*, and *full*.
- **AFLW** (Koestinger et al. 2011) consists of 24,386 faces, each with 21 landmarks. The dataset was split into 20,000 faces for training and the rest (*i.e.*, 4,386) for testing (Dong et al. 2018). Also, the left and right ears were ignored, leaving up to 19 landmarks per face sample.
- **HELEN** (Le et al. 2012) contains 2,330 images. The annotation of all 194 landmarks were used as facial prior information. We followed (Chen et al. 2018) to use the last 50 images for testing and the rest for training.
- **LFW** (Huang et al. 2007; Learned-Miller 2014) contains 13,233 face images collected from 5,750 people. Each image is labeled with the name of the person pictured. Hence, it will also be used to evaluate the recognition capabilities of super-resolved images. Note that this dataset was only used for testing.

Implementation details. We first cropped facial images about the head region, which were then resized to 128×128 . These were designated as the HR images. Then, LR images were generated by applying bicubic downsampling ($8 \times$) to the HR images, yielding a resolution of 16×16 . Then, the input LR images were reversed to match the size of the HR faces: each were up-scaled $8 \times$ using bicubic interpolation resulting in images of size 128×128 . The training images

Table 2: NMSE on 300W and AFLW. We perform the best on LR faces (bottom). Even with the proposed processing LR, while all others process HR, it still is best (top).

	300W			AFLW
	Common	Challenge	Full	
SDM (Xiong and De la Torre 2013)	5.57	15.40	7.52	5.43
LBF (Ren et al. 2014)	4.95	11.98	6.32	4.25
CFSS (Zhu et al. 2015)	4.73	9.98	5.76	3.92
MDM (Trigeorgis et al. 2016)	4.83	10.14	5.88	-
Two-stage (Lv et al. 2017)	4.36	7.56	4.99	2.17
RCSR (Wang et al. 2018)	4.01	8.58	4.90	-
CPM+SBR (Dong et al. 2018)	3.28	7.58	4.10	2.14
JASRNet (Ours)	3.20	7.44	4.03	2.03
SuperFAN (Bulat and Tzimiropoulos 2018)	5.60	10.47	6.55	3.774
FSRNet (Chen et al. 2018)	5.42	10.76	6.46	--
CPM+SBR (Dong et al. 2018)	5.42	10.65	6.45	3.87
JASRNet (Ours)	4.60	8.10	5.29	3.35

were augmented using random scaling, rotation, and horizontally flipping. Specifically, these augmentation transformations were used to make fifteen copies. Optimization was done with ADAM with a learning rate of 5.0×10^{-5} that dropped 0.5 at 20^{th} and 30^{th} epochs. The model was trained with a batch size of 8 and for a total epoch of 40 epochs. Implementation was done using PyTorch. Training took about 7 hours on Helen with a Nvidia TITAN-XP GPU.

Evaluation metrics. The metric used to evaluate landmark localization was NMSE (*i.e.*, the normalized euclidean distances between ground-truth and predicted landmarks). Following (Bulat and Tzimiropoulos 2017; Dong et al. 2018; Sagonas et al. 2013), the normalization factor is set as interocular distance for 300W and the area of the ground-truth bounding box for AFLW dataset.

For SR, we evaluated using the peak signal to noise ratio (PSNR) and structural similarity index (SSIM) (Wang et al. 2004): PSNR is computed as the mean squared error (MSE) between the SR and HR images, while SSIM accounts for the noise and edges (*i.e.*, the high-frequency content) of an image. In our experiments, we converted the RGB images to the YCbCr color space and only calculated the PSNR for the Y-channel. To focus on the face region, while ignoring the background, only the face region within the bounding box was measured when evaluating the SR images.

Comparison with state-of-the-art methods

Comparisons were made with state-of-the-art methods in both SR and face alignment. It is important to note that most existing methods only do a single task, while the proposed model does both. Furthermore, our model performs the best in both tasks. The methods that do both tasks, SuperFAN (Bulat and Tzimiropoulos 2018) and FSRNet (Chen et al. 2018), were used to compare both tasks simultaneously.

Face super-resolution results. We compared with methods used for SISR (*i.e.*, VDSR (Kim, Kwon Lee, and Mu Lee 2016), SRRes (Ledig et al. 2017), and EDSR (Lim et al. 2017)), as well as methods for face SR (*i.e.*, URDGN (Yu and Porikli 2016), TDAE (Yu and Porikli 2017), Super-

Table 3: Quantitative comparisons on LFW. Performance was measured using verification accuracy (ACC), PSNR, and SSIM. The number of parameters is also listed here.

	ACC(%)	PSNR	SSIM	Param.
HR	99.33	-	-	-
Bicubic	79.50	25.28	0.736	-
FSRNet (Chen et al. 2018)	83.75	26.63	0.800	27.14M
SuperFAN (Bulat et al. 2018)	84.08	26.83	0.808	26.41M
Ours	86.86	27.30	0.818	18.96M

FAN (Bulat and Tzimiropoulos 2018) and FSRNet (Chen et al. 2018)). For a fair comparison, we retrained aforementioned models with the same training and testing data used in the respective experiment. Qualitative comparisons clearly show that the proposed JASRNet recovers HR images with relatively more details (*i.e.*, sharper edges, more accurate facial component shapes and textures), while other methods tend to produce face images with more blur and inaccuracies (see Figure 4). Quantitative results for face SR are shown in Table 1. The proposed model achieved the highest PSNR and SSIM on 300W and HELEN dataset. Since some methods only support an upscaling factor of $\times 4$, we added an additional upscaling module ($\times 2$) to get the equivalent factor of $\times 8$. For this, we incorporated the commonly used pixel shuffle followed by a convolutional layer (Shi et al. 2016).

Face alignment results. We present face alignment results for 300W and AFLW dataset with LR image size of 16×16 and 64×64 separately. The results are summarized in Table 2. First, we compare the results of 16×16 LR images (see bottom part of Table 2). Since only a few works address the tiny face (*i.e.*, 16×16) alignment problem, we only compare the performance of proposed models with SuperFAN, FSRNet, and another state-of-the-art method CPM+SBR (Dong et al. 2018). Noticed that CPM+SBR is applied on super-resolved images using bicubic interpolation. Compared with other state-of-art methods, we show a large improvement for landmark localization on tiny faces.

Furthermore, we present results of JASRNet on faces with a resolution of 64×64 (see Table 2 (top)). Note that existing methods detect landmarks on HR (*i.e.*, 256×256) images. Still, the proposed framework is comparable for landmark localization on LR images with the others on HR.

Comparison on both tasks. To the best of our knowledge, FSRNet (Chen et al. 2018) and SuperFAN (Bulat and Tzimiropoulos 2018) were the only attempts that reported results on both tasks (*i.e.*, SR and face alignment). Thus, we compared results of both tasks with these two methods. Since one of the primary tasks for “enhancing” faces is to improve facial recognition capabilities, we also measured face verification performance on the super-resolved images. Additionally, the number of parameters used in each model is listed in Table 3. In this section, models were trained on the 300W training set, and tested on the 300W test set and the entire LFW dataset. The SR and alignment results for 300W test set are shown in Table 1 and 2, respectively. As for LFW dataset, the results for SR and facial recognition are listed in Table 3. Performance was measured using ver-

Table 4: Ablation study. To highlight the effectiveness of feature fusion and joint training.

	baseline (BL)	+feature fusion (BL_F)	joint training (JT)	+feature fusion (JT_F)	JASRNet (ours)
Super Resolution	23.41	23.50	23.55	23.58	23.69
Face Alignment	5.71	5.70	5.34	5.34	5.26

Table 5: Baseline variations of the proposed JASRNet. Trained and tested on 300W.

	Super Resolution (PSNR)	Face Alignment (NMSE)	# of Param.
Concat	23.57	5.42	–
Adding	23.69	5.29	–
One_stages	23.61	5.44	16.69M
two_stages	23.62	5.36	17.83M
Res_16	23.62	5.41	14.46M
Res_32	23.69	5.29	18.96M

ification accuracy (ACC), PSNR, and SSIM. We did not include LFW in the test for landmark localization since it does not support the 68 landmarks used as prior knowledge in all three methods. Thus, we show that our JASRNet significantly outperforms SuperFAN and FSRNet in face SR and landmark localization (see Table 1, 2, and 3). Qualitatively, the proposed method also produces more accurate landmark estimations for alignment task and much more detailed appearances and texture for SR task than the other two methods (see Figure 1, 4). Note that our model also have less parameters than SuperFAN and FSRNet (see Table 3).

Ablation study

We next measured the contributions of feature fusion, joint training, and the long skip connection. Table 4 lists the four additional variants used. Baseline (BL) only consisted of an encoder, a feature extraction module, and either a SR or alignment module. In other words, the BL omitted the feature fusion at the intermediate layers, removed the long skip connection, and was only able to handle a single task per pass (*i.e.*, either SR or face alignment, but not both). BL_F is BL with feature fusion. Joint training (JT) net was conducted by aggregating both task-specific modules to the baseline, and JT with feature fusion is JT_F. Finally, JT_F with long skip connection forms the proposed JASRNet. The training set used in this section is 300W. Note that our baseline model has even better performance while less parameters than SuperFAN (Bulat and Tzimiropoulos 2018). Reasons are three-fold: 1) batch normalization omitted in layers of residual blocks to speed up training and boost performance; 2) Pixel shuffle layers (Shi et al. 2016) used in reconstruction module instead of deconvolutional, which is used in SuperFAN; 3) Two independent modules are used in SuperFAN, *i.e.*, SR and face alignment are handled separately. This yields redundant feature maps and, hence, degrades performance.

Effects of the feature fusion. Fusing the features at the in-

termediate layers yields richer, and more efficient feature representations for SR, with BL_F and JT_F outperforming BL and JT, respectively, in SR (see Table 4). However, feature fusion has less impact on face alignment. This is because SR uses both low and high frequency information to recover HR from LR images, while landmark localization is mostly dependent on the high frequency content.

Effects of joint-task mechanism. To highlight the importance of training the two tasks jointly, we compared JT to BL and JT_F to BL_F (see Table 4). Results for both tasks (*i.e.*, SR and face alignment) show that joint-task variants (*i.e.*, JT and JT_F) significantly outperform BL and BL_F, respectfully. This validates that the joint training, in itself, contributes to the state-of-the-art performance of JASRNet.

Effects of long skip connection. The impact of the long skip connection is evident by the results: JASRNet, which is JT_F with the added skip connection, outperforms all others in both SR and landmark localization. The impact for SR stems from the skip connection forcing the network to encode sharper and more precise edges in the feature representation, as expected. However, the boosted accuracy for face alignment was less expected, yet supporting of the narrative: we believe the shared features for SR and face alignment yield additional information that complements both tasks.

Baseline variations. We also show the variations of the vanilla baseline for insights on the effects of different fusion methods (*i.e.*, concatenation vs element-wise addition), the number of residual blocks in the feature extraction module (*i.e.*, 16 vs 32), and the number of stages in the face alignment module (*i.e.*, 1 vs 2). Table 5 lists results for different settings. Clearly, element-wise addition is better for the feature fusion module in our model. Also, more residual blocks and stages improves the performance. Thus, the deeper structure and, thus, the higher capacity captures more information for the SR and face alignment tasks: as the network grows so does its potential to learn.

Conclusion

We proposed a JASRNet to exploit the maximum amount of information from tiny face images when simultaneously addressing alignment and super-resolution tasks. Extensive experiments demonstrated the proposed significantly outperforms previous state-of-the-art in SR by recovering sharper edges (*i.e.*, finer details) from HR faces. We also show large improvements for landmark localization of tiny faces (*i.e.*, 16×16). Furthermore, the proposed framework yields comparable results for landmark localization on faces of lower-resolution (*i.e.*, 64×64) to existing methods on higher-resolution (*i.e.*, 256×256).

References

- Baker, S., and Kanade, T. 2000. *Hallucinating faces*. IEEE.
- Bulat, A., and Tzimiropoulos, G. 2017. How far are we from solving the 2d & 3d face alignment problem?(and a dataset of 230,000 3d facial landmarks). In *ICCV*.
- Bulat, A., and Tzimiropoulos, G. 2018. Super-fan: Integrated facial landmark localization and sr of real-world low resolution faces in arbitrary poses with gans. In *CVPR*.
- Chen, Y.; Tai, Y.; Liu, X.; Shen, C.; and Yang, J. 2018. Fsr-net: End-to-end learning face sr with facial priors. In *CVPR*.
- Dong, X.; Yu, S.-I.; Weng, X.; Wei, S.-E.; Yang, Y.; and Sheikh, Y. 2018. Supervision-by-registration: An unsupervised approach to improve the precision of facial landmark detectors. In *CVPR*.
- Güçlütürk, Y.; Güçlü, U.; Seeliger, K.; Bosch, S.; van Lier, R.; and van Gerven, M. A. 2017. Reconstructing perceived faces from brain activations with deep adversarial neural decoding. In *NeurIPS*, 4246–4257.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*.
- Huang, G. B.; Ramesh, M.; Berg, T.; and Learned-Miller, E. 2007. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst.
- Kim, J.; Kwon Lee, J.; and Mu Lee, K. 2016. Accurate image super-resolution using very deep convolutional networks. In *CVPR*.
- Koestinger, M.; Wohlhart, P.; Roth, P. M.; and Bischof, H. 2011. Annotated facial landmarks in the wild. In *ICCVW*.
- Le, V.; Brandt, J.; Lin, Z.; Bourdev, L.; and Huang, T. S. 2012. Interactive facial feature localization. In *ECCV*.
- Learned-Miller, G. B. H. E. 2014. Labeled faces in the wild: Updates and new reporting procedures. Technical Report UM-CS-2014-003, University of Massachusetts, Amherst.
- Ledig, C.; Theis, L.; Huszár, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; et al. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*.
- Lee, C.-H.; Zhang, K.; Lee, H.-C.; Cheng, C.-W.; and Hsu, W. 2018. Attribute augmented convolutional neural network for face hallucination. In *CVPRW*.
- Lim, B.; Son, S.; Kim, H.; Nah, S.; and Mu Lee, K. 2017. Enhanced deep residual networks for single image super-resolution. In *CVPRW*.
- Liu, M.-Y.; Breuel, T.; and Kautz, J. 2017. Unsupervised image-to-image translation networks. In *NeurIPS*, 700–708.
- Liu, C.; Shum, H.-Y.; and Freeman, W. T. 2007. Face hallucination: Theory and practice. *IJCV*.
- Lv, J.-J.; Shao, X.; Xing, J.; Cheng, C.; Zhou, X.; et al. 2017. A deep regression architecture with two-stage re-initialization for high performance facial landmark detection. In *CVPR*.
- Mo, H.; Liu, L.; Zhu, W.; Yin, S.; and Wei, S. 2019. Face alignment with expression-and pose-based adaptive initialization. *IEEE Transactions on Multimedia* 21(4):943–956.
- Ranjan, R.; Patel, V. M.; and Chellappa, R. 2019. Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *TPAMI*.
- Ren, S.; Cao, X.; Wei, Y.; and Sun, J. 2014. Face alignment at 3000 fps via regressing local binary features. In *CVPR*.
- Robinson, J. P.; Li, Y.; Zhang, N.; Fu, Y.; et al. 2019. Laplace landmark localization. *arXiv preprint arXiv:1903.11633*.
- Sagonas, C.; Tzimiropoulos, G.; Zafeiriou, S.; and Pantic, M. 2013. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *ICCVW*.
- Sagonas, C.; Antonakos, E.; Tzimiropoulos, G.; Zafeiriou, S.; and Pantic, M. 2016. 300 faces in-the-wild challenge: Database and results. *Image and vision computing*.
- Shi, W.; Caballero, J.; Huszár, F.; Totz, J.; Aitken, A. P.; Bishop, R.; Rueckert, D.; and Wang, Z. 2016. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *CVPR*.
- Tai, Y.; Yang, J.; Liu, X.; and Xu, C. 2017. Memnet: A persistent memory network for image restoration. In *Proceedings of the IEEE international conference on computer vision*, 4539–4547.
- Trigeorgis, G.; Snape, P.; Nicolaou, M. A.; Antonakos, E.; and Zafeiriou, S. 2016. Mnemonic descent method: A recurrent process applied for end-to-end face alignment. In *CVPR*.
- Wang, Z.; Bovik, A. C.; Sheikh, H. R.; Simoncelli, E. P.; et al. 2004. Image quality assessment: from error visibility to structural similarity. *TIP*.
- Wang, W.; Tulyakov, S.; Sebe, N.; and err. 2018. Recurrent convolutional shape regression. *TPAMI*.
- Wei, S.-E.; Ramakrishna, V.; Kanade, T.; and Sheikh, Y. 2016. Convolutional pose machines. In *CVPR*.
- Wu, Y.; Li, J.; Kong, Y.; and Fu, Y. 2016. Deep convolutional neural network with independent softmax for large scale face recognition. In *Proceedings of the 24th ACM international conference on Multimedia*, 1063–1067. ACM.
- Xiong, X., and De la Torre, F. 2013. Supervised descent method and its applications to face alignment. In *CVPR*.
- Yu, X., and Porikli, F. 2016. Ultra-resolving face images by discriminative generative networks. In *ECCV*.
- Yu, X., and Porikli, F. 2017. Hallucinating very low-resolution unaligned and noisy face images by transformative discriminative autoencoders. In *CVPR*.
- Yu, X.; Fernando, B.; Hartley, R.; and Porikli, F. 2018. Super-resolving very low-resolution face images with supplementary attributes. In *CVPR*.
- Zhang, K.; Zhang, Z.; Li, Z.; and Qiao, Y. 2016. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*.
- Zhang, Y.; Tian, Y.; Kong, Y.; Zhong, B.; and Fu, Y. 2018. Residual dense network for image super-resolution. In *CVPR*.
- Zhu, S.; Li, C.; Change Loy, C.; and Tang, X. 2015. Face alignment by coarse-to-fine shape searching. In *CVPR*.