

Super-Identity Convolutional Neural Network for Face Hallucination

Kaipeng Zhang¹[0000-0001-6105-6532], Zhanpeng Zhang²[0000-0002-1709-4176],
Chia-Wen Cheng^{1,3}[0000-0002-5753-6530], Winston H.
Hsu^{1*}[0000-0002-3330-0638], Yu Qiao⁴[0000-0002-1889-2567], Wei
Liu⁵[0000-0002-3865-8145], and Tong Zhang⁵[0000-0002-5511-2558]

¹ National Taiwan University, Taipei, Taiwan
whsu@ntu.edu.tw

² SenseTime Group Limited, China

³ The University of Texas at Austin, Texas, USA

⁴ Shenzhen Key Lab of Computer Vision and Pattern Recognition, Shenzhen
Institutes of Advanced Technology, CAS, Shenzhen, China

⁵ Tencent AI Lab, China

Abstract. Face hallucination is a generative task to super-resolve the facial image with low resolution while human perception of face heavily relies on identity information. However, previous face hallucination approaches largely ignore facial identity recovery. This paper proposes Super-Identity Convolutional Neural Network (SICNN) to recover identity information for generating faces closed to the real identity. Specifically, we define a super-identity loss to measure the identity difference between a hallucinated face and its corresponding high-resolution face within the hypersphere identity metric space. However, directly using this loss will lead to a Dynamic Domain Divergence problem, which is caused by the large margin between the high-resolution domain and the hallucination domain. To overcome this challenge, we present a domain-integrated training approach by constructing a robust identity metric for faces from these two domains. Extensive experimental evaluations demonstrate that the proposed SICNN achieves superior visual quality over the state-of-the-art methods on a challenging task to super-resolve 12×14 faces with an $8 \times$ upscaling factor. In addition, SICNN significantly improves the recognizability of ultra-low-resolution faces.

Keywords: Face hallucination · Super Identity · Domain-Integrated Training · Convolutional Neural Networks

1 Introduction

Face hallucination, which generates high-resolution (HR) facial images from low-resolution (LR) inputs, has attracted great interests in the past few years. However, most of existing works do not take the recovery of identity information into consideration such that they cannot generate faces closed to the real identity. Fig. 1 shows some examples of hallucinated facial images generated by bicubic

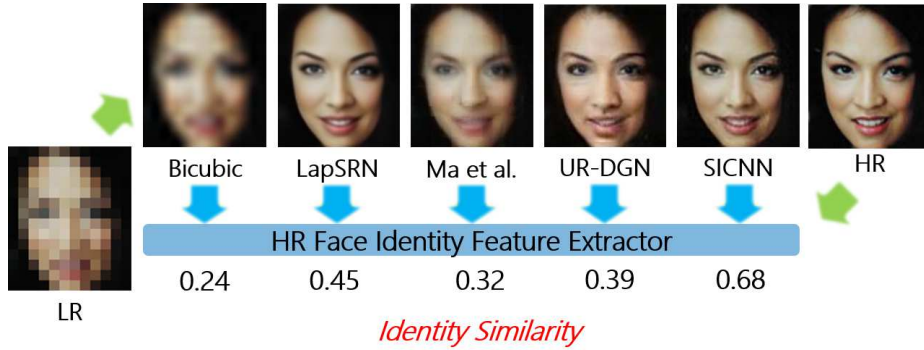


Fig. 1. Comparison of face hallucination visual quality and the performance of identity recovery over different hallucination methods. The identity similarity is computed by the cosine similarity of the identity feature.

and several state-of-the-art methods. Though they generate clearer facial images than bicubic, the identity similarities are still low, which means that they cannot recover accurate identity-related facial details. On the other hand, human perception of face heavily relies on identity information [3]. Pixel-level cues cannot fully account for the perception process of the brain. These facts suggest that recovering identity information may improve both the recognizability and performance of hallucination.

Motivated by the above observations, this paper proposes Super-Identity Convolutional Neural Network (SICNN) for identity-enhanced face hallucination. Different from previous methods, we additionally minimize the identity difference between the hallucinated face and its corresponding high-resolution face. To do so, (i) we introduce a robust identity metric space in the training process; (ii) we define a super-identity loss to measure the identity difference; (iii) we propose a novel training approach to efficiently utilize the super-identity loss. More details as follows:

For identity metric space, we use a hypersphere space [18] as the identity metric space due to its state-of-the-art performance of facial identity representation. Specifically, our SICNN is composed of a face hallucination network cascaded with a recognition network to extract identity-related feature, and an Euclidean normalization operation to project the feature into the hypersphere space.

For loss function, perceptual loss [12], computed by feature Euclidean distance, can construct convincing HR images. Differently, in our work, we need to minimize the identity distance of face pairs in the metric space. Here, we modified the perceptual loss to the super-identity loss calculated by normalized Euclidean distance (equivalent to geodesic distance) between the hallucinated face and its corresponding high-resolution face in the hypersphere identity metric space. This also facilitates our analysis on the training process (see Sec. 3.5).

For training approach, using conventional training approaches to directly train the model with super-identity loss is difficult due to the large margin between the hallucination domain and the HR domain in the hypersphere identity metric space. This is critical during the early training stage when face hallucination network cannot predict high quality hallucinated face images. Moreover, the hallucination domain keeps changing during the hallucination network learning, which makes the training with super-identity loss unstable. We summarize this challenge as a *dynamic domain divergence* problem. To overcome this problem, we propose a Domain Integrated Training algorithm that alternately updates the face recognition network and the hallucination network by minimizing the different loss in each iteration. In this alterative optimization, the hallucinated face and HR face will gradually move closer to each other in the hypersphere identity metric space while keep the discrimination of this metric space.

The main contributions of this paper are as summarized as follows:

- We propose Super-identity Convolutional Neural Network (SICNN) for enhancing the identity information in face hallucination.
- We propose Domain-Integrated Training method to overcome the problem caused by dynamic domain divergence when training SICNN.
- Compared with existing state-of-the-art hallucination methods, the SICNN achieves superior visual quality and identity recognizability when super-resolving a facial image of size 12×14 pixels with an $8 \times$ upscaling factor.

2 Related Works

Single image super-resolution (SR) aims at recovering a HR image from a LR one. Face hallucination is a kind of class-specific image SR, which exploits the statistical properties of facial images. We classify face hallucination methods into two categories: classical approaches and deep learning approach.

Classical Approaches. Subspace-based and facial components-based methods are two main kinds of classical face hallucination approaches.

For subspace-based methods. Liu et al. [17] employed a Principal Component Analysis (PCA) based global appearance model to hallucinate LR faces and a local non-parametric model to enhance the details. Ma et al. [21] used multiple local exemplar patches sampled from aligned HR facial images to hallucinate LR faces. Li et al. [16] resolved to sparse representation on local face patches. These subspace-based methods require precisely aligned reference HR and LR facial images with the same pose and facial expression.

Facial components based methods super-resolve facial parts rather than entire faces to address various poses and expressions. Tappen et al. [29] used SIFT flow to align LR images, and then deformed the reference HR images. However, the global structure is not preserved due to using local mapping. Yang et al. [33] presented a structured face hallucination method which can maintain the facial structure. However, it relies on accurate facial landmarks.

Deep Learning Approaches. Recently, deep convolutional neural networks (DCNNs) achieve remarkable progresses in a variety of computer vision tasks,

such as image classification [13], object detection [23], and face recognition [31]. Zhou et al. [37] proposed a bichannel CNN to hallucinate blurry facial images in the wild. For un-aligned faces, Zhu et al. [38] proposed to jointly learn face hallucination and facial dense spatial correspondence field estimation. The approach of [35] is a GAN-based method to generate realistic facial images. These works ignore the identity information recovery that is important for recognizability and hallucination quality. Johnson et al. [12] and Bruna et al. [2] relied on perceptual loss function closer to perceptual similarity to recover visually more convincing HR images for general image SR. In this paper we modified the perceptual loss to facilitate identity hypersphere space and propose a novel training approach to overcome the challenging while using the loss.

3 Super-Identity CNN

In this section, we will first describe the architecture of our face hallucination network. Then we will introduce the proposed super-resolution loss and super-identity loss for identity recovery. After that, we will analyze the challenge, dynamic domain divergence problem, in super-identity training. At the last, we introduce the proposed domain-integrated training algorithm to overcome this challenge.

3.1 Face Hallucination Network Architecture

As shown in Fig. 2 (a), the face hallucination network can be decomposed into feature extraction, deconvolution, mapping, and reconstruction.

We use dense block [10] to extract semantic features from LR inputs. More specifically, in the dense block, we set the growth rate to 32 and the kernel size to 3×3 . Deconvolution layer consists of learnable upscaling filters to enlarge the resolutions of input features. Mapping is implemented by a convolutional layer to reduce the dimension of features to reduce computational cost. Reconstruction also exploits a convolutional layer to predict HR images from semantic features.

Here, we denote a convolutional layer as $Conv(s, c)$ and a deconvolutional layer as $DeConv(s, c)$, where the variables s and c represent the filter size and the number of channels, respectively. In addition, PReLU [8] activation function achieves promising performance in CNN-based super-resolution [6] and we use it after each layer except the reconstruction stage.

3.2 Super-Resolution Loss

We use the pixel-wise Euclidean loss, called super-resolution loss, to constrain the overall visual appearance. For LR face input I_i^{LR} , we penalize the pixel-wise Euclidean distance between the hallucinated face and its corresponding HR face:

$$L^{SR}(I_i^{LR}, I_i^{HR}) = \|CNN_H(I_i^{LR}) - I_i^{HR}\|_2^2, \quad (1)$$

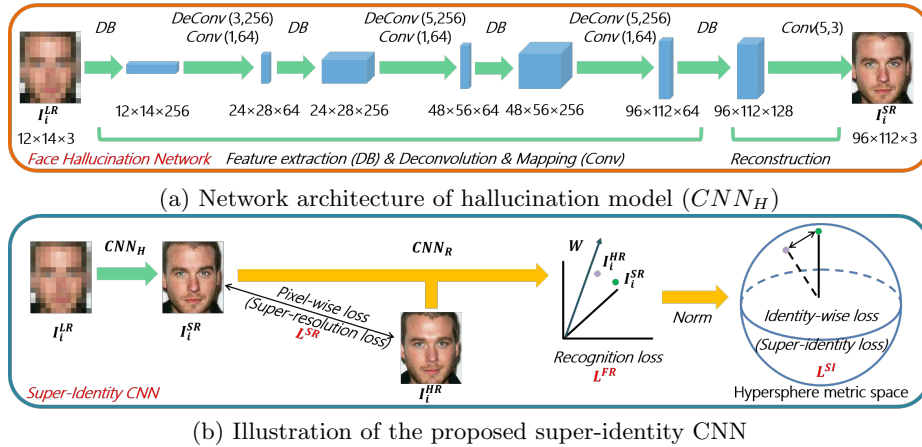


Fig. 2. Framework of our approach. (a) The network architecture of our hallucination network (CNN_H). DB denotes dense block [10]. (b) Illustration of our super-identity CNN. It uses super-resolution loss (L^{SR}), super-identity loss (L^{SI}), and recognition loss (L^{FR}) with domain-integrated training. Norm denotes Euclidean normalization, and CNN_R denotes the recognition network.

where I_i^{LR} and I_i^{HR} are the i -th LR and HR facial image pair in the training data respectively, and $CNN_H(I_i^{LR})$ represents the output of hallucination network with input I_i^{LR} . For better understanding, we also denote $CNN_H(I_i^{LR})$ as I_i^{SR} in the following text.

3.3 Hypersphere Identity Metric Space

Super-resolution loss can constrain pixel-level appearance. And we further use a constrain on the identity level. To measure the identity level difference, the first step is to find a robust identity metric space. Here we employ the hypersphere space [18] due to its state-of-the-art performance on identity representation. As shown in Fig. 2 (b), our hallucination network is cascaded with a face recognition network (i.e. CNN_R) and an Euclidean normalization operation that projects faces to the constructed hypersphere identity metric space.

CNN_R is a Resnet-like [9] CNN and more details about the network structure are introduced in our supplementary material. It is trained by A-Softmax loss function [18] which encourages the CNN to learn discriminate identity features (i.e. maximizing inter-class distance and minimizing intra-class distance) by an angular margin. In this paper, we denote this loss function as the recognition loss L^{FR} . For a face input I_i belonging to the y_i -th identity. The face recognition loss is represented as:

$$L^{FR}(I_i) = -\log\left(\frac{e^{\|CNN_R(I_i)\| \varphi(m\Theta_{y_i})}}{e^{\|CNN_R(I_i)\| \varphi(m\Theta_{y_i})} + \sum_{j \neq y_i} e^{\|CNN_R(I_i)\| \varphi(\Theta_j)}}\right), \quad (2)$$

where the Θ_{y_i} denotes the learned angle for identity y_i , $\varphi(\Theta_{y_i})$ is a monotonically decreasing function generalized from $\cos(\Theta_{y_i})$, and m is the hyper parameter of angular margin constrain. More details can be found in Sphreface [18].

3.4 Super-Identity Loss

To impose the identity information in the training process, one choice is to use a loss computed by features Euclidean distance between face pairs, such as perceptual loss [12]. However, in this paper, since our goal is to minimize identity distance in hypersphere metric space, the original perceptual loss, computed by L2 distance is not the best choice in our task. Therefore, we propose a modified perceptual loss, called Super-Identity (SI) loss, to compute the normalized Euclidean distance (equivalent to geodesic distance). This modification makes the loss directly related to identity in hypersphere space and facilitate our investigation in Sec. 3.5.

For a LR face input I_i^{LR} , we penalize the normalized Euclidean distance between the hallucinated face and its corresponding HR face in the constructed hypersphere identity metric space:

$$L^{SI}(I_i^{LR}, I_i^{HR}) = \left\| \widehat{CNN_R(I_i^{SR})} - \widehat{CNN_R(I_i^{HR})} \right\|_2^2 \quad (3)$$

where $CNN_R(I_i^{SR})$ and $CNN_R(I_i^{HR})$ are the identity features extracted from face recognition model (CNN_R) for facial images I_i^{SR} and I_i^{HR} , respectively. $\widehat{CNN_R(I_i^{SR})} = \frac{CNN_R(I_i^{SR})}{\|CNN_R(I_i^{SR})\|_2}$ is the identity representation projected to the unit hypersphere.

In addition to L^{SI} , we want to have some discussions about perceptual loss beyond our work. In general, the perceptual loss is computed by L2 distance. However, in most CNNs, inner-product operation is used in fully-connected and convolutional layers. These outputs are related to the feature’s norm, weight’s norm and the angular between them. Therefore, for different tasks and different metric space (e.g. [19, 5, 24]), some modifications about computational metric space of perceptual loss are necessary (L^{SI} is one of the cases).

3.5 Challenges of Training with Super-Identity Loss

Super-identity loss imposes an identity level constrain. We examine different training methods as follows:

Baseline training approach I. A straightforward way to train our framework is jointly using the L^{SR} , L^{SI} and L^{FR} to train both CNN_H and CNN_R from scratch. The optimization objective can be represented as:

$$\min_{\theta_{CNN_H} \theta_{CNN_R}} \frac{1}{n} \sum_{i=1}^n L^{SR}(I_i^{LR}, I_i^{HR}) + \alpha L^{SI}(I_i^{LR}, I_i^{HR}) + \beta L^{FR}(I_i^{SR}, I_i^{HR}), \quad (4)$$

where α and β denotes the loss weight of the L^{SI} and L^{FR} respectively, θ_{CNN_H} and θ_{CNN_R} denotes the learnable parameters.




Fig. 3. Face hallucination examples produced by CNN_H trained by different training approaches. These four columns of results are produced by baseline training approach I, II, III and the proposed domain-integrated training approach respectively. It is clear that our approach achieves the best result while other results are noisy. This figure is best viewed in color. Please *zoom in* for better comparison.

Observation I. This training approach generates artifacts (see Fig. 3, first column) and the loss is too difficult to converge. The reasons may come from: (1) In the early training stage, the hallucinated faces are quite different from HR faces, so the CNN_R is too difficult to be optimized from scratch. (2) The objective of L^{FR} (i.e. minimizing the intra-class variance) is different from the objective of L^{SI} and L^{SR} loss (minimizing the pair-wise distance), which is disadvantageous to CNN_R and CNN_H learning. So, we cannot use the L^{SI} in CNN_R learning and also cannot use the L^{FR} in CNN_H learning.

Baseline training approach II. To solve above problems, one possible training approach used in perceptual loss [12] can be used. In particular, we train a CNN_R using HR faces and then jointly use the L^{SR} and the L^{SI} to train the CNN_H . The joint objective of L^{SI} and L^{SR} can be represented as:

$$\min_{\theta_{CNN_H}} \frac{1}{n} \sum_{i=1}^n L^{SR}(I_i^{LR}, I_i^{HR}) + \alpha L^{SI}(I_i^{LR}, I_i^{HR}), \quad (5)$$

Observation II. We have two observations while using this training approach: (1) The L^{SI} is difficult to converge. (2) The visual results are noisy (see Fig. 3, second column). To investigate these challenges, we first visualized the learned identity features (after Euclidean normalization, as shown in Fig. 4) and found that there exists a large margin between the hallucination domain and the HR domain. We formulate this challenge as domain divergence problem. It specifies the failure of the CNN_R , trained by HR faces, to project faces from hallucination domains to a measurable hypersphere identity metric space. In other words, this face recognition model cannot extract effective identity representation for hallucinated faces. This makes the L^{SI} very difficult to converge and easily get stuck in local minima (i.e. occur many noises in hallucination results).

Baseline training approach III. To overcome the domain divergence challenge, a straightforward alternately training strategy can be used. In particular,




Fig. 4. The distribution of identity features (after Euclidean normalization) from hallucination domain (triangle) and HR domain (dot). These identities are randomly selected from the training set. Different colors denote different identities. We use t-SNE [30] to reduce the dimensions for better understanding. We can observe that there is a large gap between above two domains in the identity metric space.

we first trained a CNN_H only using the L^{SR} . Then we trained a CNN_R using hallucinated faces and HR faces. Finally, we finetune the CNN_R jointly using the L^{SR} and the L^{SI} following baseline training approach II.

Observation III. Although this alternately training strategy seems able to overcome the domain divergence problem, it still produces artifacts (as shown in Fig. 3, third column). The reason is that the hallucination domain keeps changing when the CNN_H is being updated. If the hallucination domain has changed, the face recognition model cannot extract effective and measurable identity representation of hallucinated faces anymore.

In short, above observations can be concluded into a dynamic domain divergence problem as following: a large margin exists between the hallucination domain and HR domain and the hallucination domain keeps changing if the hallucination model keeps learning.

3.6 Domain-Integrated Training Algorithm

To overcome the dynamic domain divergence problem, we propose a new training procedure. From above the above observations, we see that alternately training strategy (Baseline Training Approach III) can alleviate the dynamic domain divergence problem. We further propose to do this alternately training in each iteration.

Algorithm 1 Mini-batch SGD based domain-integrated training approach

Input: Face recognition model CNN_R trained by HR facial images, face hallucination model CNN_H trained by L^{SR} , minibatch size N , LR and HR facial image pairs $\{I_i^{LR}, I_i^{HR}\}$.

Output: SICNN.

- 1: **while** not converge **do**
- 2: Choose one minibatch of N LR and HR image pairs $\{I_i^{LR}, I_i^{HR}\}, i = 1, \dots, N$.
- 3: Generate one minibatch of N hallucinated facial images I_i^{SR} from $I_i^{LR}, i = 1, \dots, N$, where $I_i^{SR} = CNN_H(I_i^{LR})$.
- 4: Update the recognition model CNN_R by descending its stochastic gradient:

$$\nabla_{\theta_{CNN_R}} \frac{1}{N} \sum_{i=1}^N L^{FR}(\{I_i^{SR}, I_i^{HR}\})$$
- 5: Update the hallucination model CNN_H by descending its stochastic gradient:

$$\nabla_{\theta_{CNN_H}} \frac{1}{N} \sum_{i=1}^N L^{SR}(I_i^{LR}, I_i^{HR}) + \alpha L^{SI}(I_i^{LR}, I_i^{HR})$$
- 6: **end while**

More specifically, we first train a CNN_R using HR facial images and a CNN_H using the L^{SR} . Then, we propose to use domain-integrated training approach (Algorithm 1) to finetune CNN_R and CNN_H alternately in each iteration.

In particular, in each iteration, we first update the CNN_R using the recognition loss, which allows the CNN_R to perform accurate identity representation in this mini-batch of faces from different domains. Then, we jointly use the L^{SR} and the L^{SI} to update the CNN_H . This training approach can encourage the CNN_R to construct a robust mapping from faces to the measurable hypersphere identity metric space in each iteration for L^{SI} optimization whatever the CNN_H is changing. The alternative optimization process is conducted until converged. Some hallucination examples are shown in Fig. 3, fourth column, where we can observe a much better visual result with this training approach.

3.7 Comparison to Adversarial Training

Domain-Integrated (DI) training and adversarial training [7] can be related to their alternative learning strategy. But they are quite different in several aspects as follows:

(1) Generally speaking, DI training is essentially a cooperative process in which CNN_H collaborates with CNN_R to minimize the identity difference. The learning objective is the same in each sub-iteration. However, in adversarial training, generator and discriminator compete against each other to improve the performance. The learning objective is alternatively challenging during two models learning.

(2) The loss functions and optimization style are different. In DI training, we minimize L^{FR} in CNN_R constructing a marginal identity metric space and then minimize L^{SI} for CNN_H reducing pair-wise identity difference. Differently, in adversarial training, the classification loss is minimized for discriminator learning and maximized for generator learning.

4 Experiments

In this section, we will first describe the training and testing details. Then we perform an ablation study to evaluate the effectiveness of the proposed Super-Identity loss and Domain-Integrated training. Further, we evaluate our proposed method with other state-of-the-art methods. At the last, we evaluate the benefit of our method for low-resolution face recognition. *More evaluations are included in our supplementary material.*

4.1 Training Details

Training data. For a fair comparison with other state-of-the-art methods, we do face alignment in facial images. In particular, we use similarity transformation based on five landmarks detected by MTCNN [36]. *We have removed the images and identities overlap between training and testing.*

For face recognition training, we use web-collected facial images including CASIA-WebFace [34], CACD2000 [4], CelebA [20], VGG Faces [22] as *Set A*. It roughly goes to 1.5M images of 17,680 unique persons.

For face hallucination training, we select 1.1M HR facial images (larger than 96×112 pixels) from the same 1.5M images as *Set B*.

Training details. For recognition model training, we use *Set A* with the batch size of 512 and m (angular margin constrain in Eq. 2) of 4. The learning rate is started from 0.1 and divided by 10 at the 20K, 30K iterations. The training process is finished at 35K iterations.

For hallucination model training, we use *Set B* with the batch size of 128. The learning rate is started from 0.02 and divided by 10 at the 30K, 60K iterations. A complete training is finished at 80K iterations.

For domain-integrated training, we use *Set B* with the batch size of 128 for CNN_H and 256 for CNN_R . The learning rate is started from 0.01 and divided by 10 at the 6K iterations. A complete training is finished at 9K iterations.

4.2 Testing Details

Testing data. We randomly select 1,000 identities with 10,000 HR facial images (larger than 96×112 pixels) from UMD-Face [1] dataset as *Set C*. The dataset is used for face hallucination and identity recovery evaluation.

Evaluation protocols. In this section, we perform three kinds of evaluations: (1) Visual quality. (2) Identity recovery. (3) Identity recognizability. For visual quality evaluation, we report several visual examples results on *Set C*.

For identity recovery, we evaluate the performance of recovering identity information while super-resolving faces. In particular, we use the CNN_R trained by *Set A* as identity features extractor. And the identity features are taken from the output of the first fully connected layer. Then we compute the identity similarity (i.e. cosine similarity) between the hallucinated face and its corresponding HR faces on *Set C*. The average similarities over the testing set are reported.

For identity recognizability, we evaluate the recognizability of hallucinated faces. In particular, we first downsample *Set A* to 12×14 pixels as *Set A - LR*. Then we use different methods to super-resolve *Set A - LR* to 96×112 pixels as different *Set A - SR*. At last, we use the *Set A - SR* to train different CNN_R and evaluate them on LFW [11] and YTF [32].

4.3 Ablation Experiment

Loss weight. The hyper parameter α (see Algorithm 1) dominates the identity recovery. To verify the effectiveness of the proposed Super-Identity loss, we vary α from 0 (i.e. only use super-resolution loss) to 32 to learn different models. From Tab. 1 and Fig. 5, we observe that larger α make the facial images sharper with more details and brings the better performance of identity recovery and recognizability. But too large α also makes the texture look slightly unnatural. And, since the performances of identity recovery and identity recognizability are stable when α is larger than 8, we fix α to 8 in other experiments.

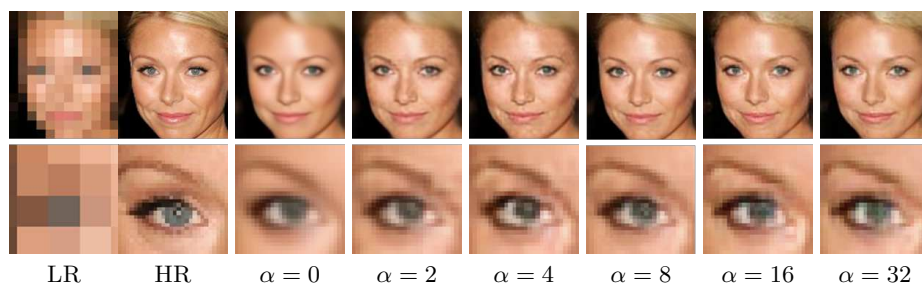


Fig. 5. Face hallucination examples generated by models trained with different loss weight α . It is clear that choosing larger α can make the facial images sharper with more details. Please *zoom in* for better comparison.

α	0	2	4	8	16	32
Identity Similarity	0.4418	0.5134	0.5639	0.5978	0.6041	0.6101
LFW Accuracy	97.61%	97.88%	98.05%	98.25%	98.23%	98.16%
YTF Accuracy	93.20%	93.48%	93.56%	93.82%	93.84%	93.76%

Table 1. Quantitative comparison of different α on identity recovery and identity recognizability evaluation. Larger α brings better performance and it is stable when α is larger than 8.

Training approach. We evaluate different training approaches introduced in Sec. 3.5 and Sec. 3.6. Some visual results are shown in Fig. 3. We can see that

Training Approach	I	II	III	Domain-Integrated Training
Identity Similarity	0.3875	0.4829	0.5132	0.5978
LFW Accuracy	97.16%	97.46%	97.58%	98.25%
YTF Accuracy	92.98%	93.32%	93.34%	93.84%%

Table 2. Quantitative comparison of different training approaches on identity recovery and identity recognizability evaluation. The results demonstrate the superiority of our proposed domain-integrated training.

Domain-Integrated training achieves the best visual results. Besides, from Tab. 2, Domain-Integrated training also achieves the best performance of identity recovery and identity recognizability.

4.4 Evaluation on Face Hallucination

We compare SICNN with other state-of-the-art methods and bicubic interpolation on *Set C* for face hallucination. In particular, we follow EnhanceNet [25] training another UR-DGN, called UR-DGN*, with additional perceptual loss computed in end of the second and the last ResBlock in CNN_R . *All methods are re-trained in same training set - Set B.*

Some visual examples are shown in Fig. 6. More visual results are included in our supplementary material. We also report the results of average Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity (SSIM) in Tab. 3. But as the claim of other works [12, 25, 15], PSNR and SSIM results are useless for semantic super-resolution evaluation while visual quality and recognizability are more valuable.

From the visual results, it is clear that our method achieves the best results over other methods. We analyze the results as follows:

(1) For Ma *et al.*'s method, exemplar patches based, the results are over-smooth and suffer from obvious blocking for such low low-resolution input with large up-sampling scale.

(2) For LapSRN [14], since it is based on L2 pixel-wise loss, it makes the hallucinated faces over-smooth.

(3) For UR-DGN [35], it jointly uses pixel-wise Euclidean loss and adversarial loss to generate a realistic facial image closest to the average of all potential images. Thus, though the generated facial images look realistic, they are quite different from the original HR images.

(4) For UR-DGN*, it uses an additional loss - perceptual loss computed in our CNN_R as the pair-wise semantic loss for identity recovery. Though this pixels-wise loss + adversarial loss + perceptual loss is the state-of-the-art super-resolution training approach (i.e. EnhancementNet[25]). It still achieves inferior results than ours.




Fig. 6. Comparison with the state-of-the-art methods on hallucination test dataset. It is clear that our method achieves the best hallucination visual quality. Please *zoom in* for better comparison. More visual results are included in our supplementary material.

Method	Bicubic	Ma <i>et al.</i>	LapSRN	UR-DGN	UR-DGN*	SICNN
PSNR (db)	23.1323	23.8606	26.1451	24.1857	25.2859	26.8945
SSIM	0.6093	0.6571	0.7417	0.6764	0.7224	0.7689

Table 3. Quantitative hallucination comparison of different training approaches.

4.5 Evaluation on Identity Recovery

We perform an evaluation on identity recovery with other state-of-the-art methods. All models for evaluation are the same as last experiment (i.e. Sec. 4.4).

From the Tab. 4, we observe that our method achieves the best performance. Besides, we also observe that UR-DGN, trained by pixels-wise loss and adversarial loss, even shows inferior performance than LapSRN though with sharper visual results (See Sec. 4.4). It means that UR-DGN will lose some identity information while super-resolving a face because the adversarial loss is not a pair-wise loss. And if add perceptual loss (i.e. UR-DGN*), pair-wise semantic loss, the results can be improved, but still inferior to our method.

4.6 Evaluation on Identity Recognizability

Follow last two experiments (i.e. Sec. 4.4, 4.5)., we further perform an evaluation on identity recognizability with other state-of-the-art methods.

From the Tab. 4, we observe that our method achieves the best performance. We also obtain similar observations as last experiment. Besides, we also observe that though several methods (LapSRN, Ma *et al.*, and UR-DGN) obtain better visual results than Bicubic interpolation, the identity recognizability of super-resolved face is similar or even inferior. It means that these methods cannot generate discriminative faces with better identity recognizability.

Method	Bicubic	Ma <i>et al.</i>	LapSRN	UR-DGN	UR-DGN*	SICNN
Identity Similarity	0.2913	0.3823	0.4361	0.3682	0.5267	0.5978
LFW Acc.	97.51%	97.58%	97.46%	97.20%	98.01%	98.25%
YTF Acc.	93.08%	93.26%	93.10%	92.78%	93.54%	93.82%

Table 4. Quantitative comparison on identity recovery and identity recognizability evaluation. The results demonstrate the superiority of our proposed method.

Method	Ours	CNN_R	Human	[28]	[27]	[26]	[22]	[31]	[18]
Input Size	12×14	96×112	Original	152×152	47×55	224×224	224×224	96×112	96×112
LFW Acc.	98.25%	99.48%	97.53%	97.35%	98.70%	99.63%	98.95%	99.28%	99.42%
YTF Acc.	93.82%	95.38%	-	91.4%	93.2%	95.1%	97.3%	94.9%	95.0%

Table 5. Face verification performance of different methods on LFW [11] and YTF [32] benchmark. It shows that our method can help the recognition model to archive high accuracy with ultra-low-resolution inputs.

4.7 Evaluation on Low-Resolution Face Recognition

To evaluate the benefit of our method for low-resolution face recognition, we compare our method ($SICNN + CNN_R$) with other state-of-the-art recognition methods on LFW [11] and YTF [32] benchmark.

From the results in Table 5, we find that these methods’ input sizes are relatively large (area size from $15.3\times$ to $298\times$ compared with our method). Moreover, using our face hallucination method, the recognition model can still achieve reasonable results in such ultra-low resolution. We also tried using un-aligned faces in training and testing and our proposed method still can achieve similar improvement of performance.

5 Conclusion

In this paper, we present Super-Identity CNN (SICNN) to enhance the identity information during super resolving face images of size 12×14 pixels with an $8\times$ upscaling factor. Specifically, SICNN aims to minimize the identity difference between the hallucinated face and its corresponding HR face. In addition, we propose a domain-integrated training approach to overcome the dynamic domain divergence problem when training SICNN. Extensive experiments demonstrate that SICNN not only achieves superior hallucination results but also significantly improves the performance of low-resolution face recognition.

6 Acknowledgement

This work was supported in part by MediaTek Inc and the Ministry of Science and Technology, Taiwan, under Grant MOST 107-2634-F-002 -007. We also benefit from the grants from NVIDIA and the NVIDIA DGX-1 AI Supercomputer.

References

1. Bansal, A., Nanduri, A., Castillo, C., Ranjan, R., Chellappa, R.: Umdfaces: An annotated face dataset for training deep networks. arXiv:1611.01484 (2016)
2. Bruna, J., Sprechmann, P., LeCun, Y.: Super-resolution with deep convolutional sufficient statistics. ICLR (2016)
3. Chang, L., Tsao, D.Y.: The code for facial identity in the primate brain. *Cell* **169**(6), 1013–1028 (2017)
4. Chen, B.C., Chen, C.S., Hsu, W.H.: Face recognition and retrieval using cross-age reference coding with cross-age celebrity dataset. *TMM* **17**(6), 804–815 (2015)
5. Chunjie, L., Qiang, Y., et al.: Cosine normalization: Using cosine similarity instead of dot product in neural networks. arXiv (2017)
6. Dong, C., Loy, C.C., Tang, X.: Accelerating the super-resolution convolutional neural network. In: ECCV. pp. 391–407 (2016)
7. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: NIPS. pp. 2672–2680 (2014)
8. He, K., Zhang, X., Ren, S., Sun, J.: Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: ICCV. pp. 1026–1034 (2015)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR. pp. 770–778 (2016)
10. Huang, G., Liu, Z., Weinberger, r.K., Maaten, L.v.d.: Densely connected convolutional networks (2017)
11. Huang, G.B., Learned-Miller, E.: Labeled faces in the wild: Updates and new reporting procedures. Dept. Comput. Sci., Univ. Massachusetts Amherst, Amherst, MA, USA, Tech. Rep pp. 14–003 (2014)
12. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: ECCV. pp. 694–711 (2016)
13. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS. pp. 1097–1105 (2012)
14. Lai, W.S., Huang, J.B., Ahuja, N., Yang, M.H.: Deep laplacian pyramid networks for fast and accurate super-resolution. CVPR (2017)
15. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network (2017)
16. Li, Y., Cai, C., Qiu, G., Lam, K.M.: Face hallucination based on sparse local-pixel structure. *PR* **47**(3), 1261–1270 (2014)
17. Liu, C., Shum, H.Y., Freeman, W.T.: Face hallucination: Theory and practice. *IJCV* **75**(1), 115–134 (2007)
18. Liu, W., Wen, Y., Yu, Z., Li, M., Raj, B., Song, L.: Spheraface: Deep hypersphere embedding for face recognition (2017)
19. Liu, W., Zhang, Y.M., Li, X., Yu, Z., Dai, B., Zhao, T., Song, L.: Deep hyperspherical learning. In: NIPS. pp. 3953–3963 (2017)
20. Liu, Z., Luo, P., Wang, X., Tang, X.: Deep learning face attributes in the wild. In: ICCV. pp. 3730–3738 (2015)
21. Ma, X., Zhang, J., Qi, C.: Hallucinating face by position-patch. *PR* **43**(6), 2224–2236 (2010)
22. Parkhi, O.M., Vedaldi, A., Zisserman, A.: Deep face recognition. In: BMVC. vol. 1, p. 6 (2015)

23. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: NIPS. pp. 91–99 (2015)
24. Rippel, O., Paluri, M., Dollar, P., Bourdev, L.: Metric learning with adaptive density discrimination. ICLR (2016)
25. Sajjadi, M.S., Scholkopf, B., Hirsch, M.: Enhancenet: Single image super-resolution through automated texture synthesis. In: ICCV. pp. 4491–4500 (2017)
26. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: CVPR. pp. 815–823 (2015)
27. Sun, Y., Wang, X., Tang, X.: Deeply learned face representations are sparse, selective, and robust. In: CVPR. pp. 2892–2900 (2015)
28. Taigman, Y., Yang, M., Ranzato, M., Wolf, L.: Deepface: Closing the gap to human-level performance in face verification. In: CVPR. pp. 1701–1708 (2014)
29. Tappen, M.F., Liu, C.: A bayesian approach to alignment-based image hallucination. In: ECCV. pp. 236–249 (2012)
30. Van Der Maaten, L.: Accelerating t-sne using tree-based algorithms. JMLR **15**(1), 3221–3245 (2014)
31. Wen, Y., Zhang, K., Li, Z., Qiao, Y.: A discriminative feature learning approach for deep face recognition. In: ECCV. pp. 499–515 (2016)
32. Wolf, L., Hassner, T., Maoz, I.: Face recognition in unconstrained videos with matched background similarity. In: CVPR. pp. 529–534 (2011)
33. Yang, C.Y., Liu, S., Yang, M.H.: Structured face hallucination. In: ECVP. pp. 1099–1106 (2013)
34. Yi, D., Lei, Z., Liao, S., Li, S.Z.: Learning face representation from scratch. arXiv:1411.7923 (2014)
35. Yu, X., Porikli, F.: Ultra-resolving face images by discriminative generative networks. In: ECCV. pp. 318–333 (2016)
36. Zhang, K., Zhang, Z., Li, Z., Qiao, Y.: Joint face detection and alignment using multitask cascaded convolutional networks. SPL **23**(10), 1499–1503 (2016)
37. Zhou, E., Fan, H., Cao, Z., Jiang, Y., Yin, Q.: Learning face hallucination in the wild. In: AAAI. pp. 3871–3877 (2015)
38. Zhu, S., Liu, S., Loy, C.C., Tang, X.: Deep cascaded bi-network for face hallucination. In: ECCV. pp. 614–630 (2016)