

QGFace: Quality-Guided Joint Training For Mixed-Quality Face Recognition

Youzhe Song, and Feng Wang*

摘要—

THE quality of a face crop in an image is decided by many factors such as camera resolution, distance, and illumination condition. This makes the discrimination of face images with different qualities a challenging problem in realistic applications. However, most existing approaches are designed specifically for high-quality (HQ) or low-quality (LQ) images, and the performances would degrade for the mixed-quality images. Besides, many methods ask for pre-trained feature extractors or other auxiliary structures to support the training and the evaluation. In this paper, we point out that the key to better understand both the HQ and the LQ images simultaneously is to apply different learning methods according to their qualities. We propose a novel quality-guided joint training approach for mixed-quality face recognition, which could simultaneously learn the images of different qualities with a single encoder. Based on quality partition, classification-based method is employed for HQ data learning. Meanwhile, for the LQ images which lack identity information, we learn them with self-supervised image-image contrastive learning. To effectively catch up the model update and improve the discriminability of contrastive learning in our joint training scenario, we further propose a proxy-updated real-time queue to compose the contrastive pairs with features from the genuine encoder. Experiments on the low-quality datasets SCface and Tinyface, the mixed-quality dataset IJB-B, and five high-quality datasets demonstrate the effectiveness of our proposed approach in recognizing face images of different qualities.

Index Terms—Face recognition, discriminative representations, contrastive regularization, mixed-quality face recognition.

Youzhe Song and Feng Wang are with Shanghai Key Laboratory of Multidimensional Information Processing, School of Computer Science and Technology, East China Normal University, Shanghai, 200062 China.

*Corresponding author (Email:fwang@cs.ecnu.edu.cn).

I. INTRODUCTION

Existing face recognition (FR) methods have achieved near-perfect performances on high-quality (HQ) face images [1]–[3]. However, their performances significantly degrade when being applied to low-quality (LQ) images. Some methods aim to improve performance in specific image domains, such as light field images [4] and near-infrared images [5] and hard to be applied to general LQ images. Some other works focus on the face recognition where LQ images are adopted for the matching between the given probe images and the gallery images [6]–[11]. However, due to the domain gap between the HQ and the LQ images, the discriminability for HQ images is somewhat reduced. It is hard to effectively learn from both the LQ and the HQ images simultaneously. In reality, the quality of a face image used for recognition is influenced by numerous factors such as camera resolution, distance, occlusion, and illumination condition. For instance, a high-resolution smart phone camera could produce low-quality face crops when there are too many faces in the picture or the photographer focuses on the scenery in the background. Meanwhile, in the surveillance situations, the image qualities also vary when the distance is changing (the setting in SCface [12] dataset). Domain gaps also exist among these low-quality images. In some limited situations where the image qualities are fixed, it is possible to select the proper model according to the image qualities. However, for the applications such as photo albums and complicated surveillance systems with various cameras, the qualities of face images vary a lot. Existing methods only focusing on either HQ or LQ images are not suitable for such scenarios.

The problems with the mixed-quality face recognition

using a single feature extractor lie in the following aspects. First, due to the domain gap between the HQ and the LQ data, it is hard to learn from and inference with images of different qualities simultaneously. Most existing methods emphasize either the HQ or the LQ images, while miss the other one. Some quality-invariant methods [9]–[11], [13], [14] attempt to understand the LQ images with the guidance of a pretrained HQ model. This explicitly teaches the model to lean to the LQ images and thus the performance degrades for the HQ images. Besides, as the HQ images are encoded by a fixed model, the compromise between images of difference qualities is also under-optimized. Second, we argue that there is a semantic problem when applying classification-based method on the LQ images. The identities of faces are hard to predict due to the natural variations of their poses, ages, and makeups for the HQ images. Compared with the HQ images, the LQ images typically have fewer details, which makes it challenging for the model to accurately classify them. This confusion arises because the model may struggle to differentiate between natural variations in the images and disturbances caused by the lower image qualities. As a result, the LQ images usually stumble the performance of the feature extractor trained with the classification loss [15].

To address the above problems, we propose a novel approach namely quality-guided joint training for mixed-quality face recognition (QGFace). Our goal is to learn both the HQ and the LQ face images simultaneously, and recognize face images without distinction of qualities with a single feature extractor. To effectively utilize the data of different qualities, we design a quality partitioning strategy and apply different supervision signals according to the qualities of the images. Specifically, for the HQ images, the classification-based method is applied as they have achieved near-perfect performance for face recognition [1], [2], [15], [16]. For the LQ images which could mislead the classification loss, we adopt self-supervised contrastive learning which learns the face images at the instance level and alleviates the requirements for the image details. By purposefully taking advantages of the classification methods and the contrastive learning, we get an encoder which is applicable for images with different qualities.

Furthermore, with the joint training strategy, we only need a single feature extractor during training and evaluation. This makes our method highly extensible, which could serve as the base component for future works.

Nevertheless, taking contrastive learning from the unsupervised learning into the joint training scenario is non-trivial. Face identification asks for searching the related identity of a given image. Classification-based methods contain thousands of image-identity pairs to learn discriminative features which could meet the above requirement. For contrastive learning, it usually takes a large number of sample pairs to improve the stability and the accuracy. Momentum queue is a common choice [17]. It generates representations by an extra encoder which is momentum-updated with the training encoder. However, it loses track of the model which is also guided by the classification supervision signal, and the produced features are whitened as the encoder keeps average of all the past model statuses. Inspired by BroadFace [18], we propose a novel proxy-updated real-time contrastive queue to supply positive pairs from the genuine model with a large number of negative samples which is dynamically updated. The proxy is represented by the corresponding weights of an identity in the classifier.

The contributions of this paper are summarized as follows:

- We propose a quality-guided joint training strategy for face recognition (QGFace) on the mixed-quality images. To better understand the LQ images which could mislead the classification methods, we partition the images into the HQ and the LQ subsets. Classification loss and contrastive loss are then applied on the two subsets respectively. In this way, we can comprehensively learn from both the LQ and the HQ faces to train a single encoder.
- To support the feature search in face identification, we propose a proxy-updated real-time contrastive queue to ensure the discriminability of the LQ images learned by contrastive learning. The identity proxy representations are taken to update the queue features encoded in the previous training step.
- We conduct extensive experiments on high-quality,

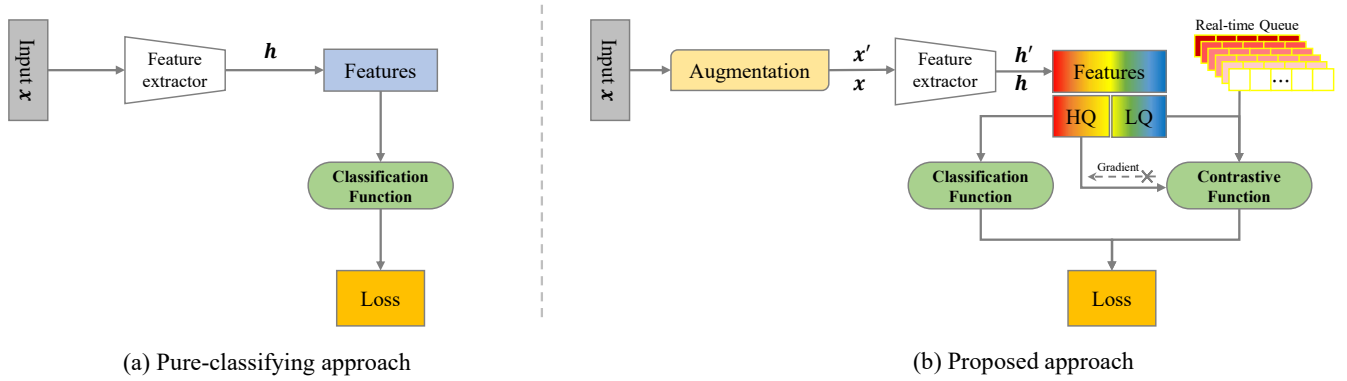


图 1. Conventional pure classification method vs. our QGFace. (a) An FR training pipeline with a single classification loss. All the input images are taken into a single main loss. (b) Our pipeline additionally takes contrastive loss for low-quality images. We apply data augmentation on images and split their features into two parts: HQ and LQ ones. The HQ images are supervised by the classification loss, while the LQ images and their related HQ images are sent into the contrastive function. The gradient flow of HQ data is stopped in contrastive learning. A real-time queue is designed to provide an effective feature queue and support the large-scale feature comparison.

low-quality, and mixed-quality datasets. Our approach achieves competitive results on different categories of datasets compared with the SOTA methods, which demonstrates the effectiveness of our QGFace.

II. RELATED WORKS

A. Encoder-only methods

Data augmentation is often employed to produce LQ images since they can provide similar domain information as the LQ evaluation datasets [9]–[11]. However, these on-the-fly augmentation methods usually degrade the discriminability of the images and make them hard to be identified in model training. Recent works discover the relationship between the feature norm and the image quality [15], [19]. With the awareness of image qualities, the loss functions are designed to assign different importance to the images according to their qualities. They attempt to improve the performance on the LQ images while keeping the high accuracy on the HQ images. Nevertheless, the incompatibility between the LQ data and the classification training process prohibits the encoder from effectively learning the hard LQ images. PENB [20] forces the alignment between the distribution of LQ and HQ data in a global view, which may suffer from such a problem. Instead of learning LQ data with classification-based method, we introduce contrastive learning to better understand them with the instance-level supervision. Furthermore, by quality partitioning, we can

prevent the LQ images from misleading the classification-based HQ learning.

B. Quality-invariant Methods

To deal with the domain gap between the LQ and the HQ face images, some methods attempt to transfer the knowledge from the teacher model trained with the HQ images to the student model which focuses on the LQ images. In [10], [11], a pretrained teacher network is frozen and then the student network is trained to extract features similar to the teacher with the down-sampled HQ images. TCN [9] takes a contrastive loss to constrain the relationship between the HQ and the LQ images. RAN [13] proposes multi-resolution Generative Adversarial Networks (GAN) to synthesize the cross-resolution images and a feature adaption network to enhance the communication between the feature extractors for the HQ and the LQ images. DDL [14] splits the data into two subsets according to the predefined semantic difficulty including yaw, race, and distance during image capture. A similarity-based KL-divergence loss is applied on the two subsets.

These methods pursue the knowledge transfer from the HQ to the LQ images. However, with such a single-way learning process, the performances of the learned models on the HQ images are stumbled. Our QGFace partitions the training data according to their qualities and applies different supervising methods on them, which not only suits

for the LQ face recognition, but also keeps the ability of recognizing the HQ images.

C. Synthesis-based Methods

To generalize the high performance of existing high-quality face recognition to different qualities, two straightforward ways are to super-resolve the LQ images for recognition [6]–[8], [21] and to synthesize meaningful multi-quality training data. As they can only learn the mapping between the artificially down-sampled images and the original one, the training of generation model usually needs a well-trained feature extractor. In [6]–[8], [21]–[25], a pretrained feature extractor is taken to supervise the identities of the synthetic images. In [26], a GAN is trained to learn how to downgrade the high-resolution images to generate realistic low-resolution images. In [27], a DCR model is taken to reserve the identity information, while in [28], [29] the network structure is modified to improve the generation quality. Owing to the absence of large-scale labeled LR face datasets, the approaches in [30], [31] synthesize low-quality images from the high-quality dataset and make it possible to train a mixed-quality FR model. [24] only focuses on a specific type of LQ images which is the partial face images. [25] reconstructs the HQ face images with singular value decomposition to improve the reliability.

These methods could obtain satisfactory HQ images to mitigate the effect of resolution difference or lead the model to learn to discriminate LQ face images. However, as an extra generative module is needed, the overall framework is not end-to-end and computationally expensive. Besides, a well-trained feature extractor is still needed in these methods, which highly dominates the training of the synthesis-based model and the performance of the recognition process. By directly training a single feature extractor which is applicable to both the HQ and the LQ data in an end-to-end way, our method is more elegant and orthogonal with the synthesis-based methods.

III. METHOD

Fig. 1(b) illustrates the framework of our quality-guided joint training for mixed-quality face recognition.

We first partition the features into HQ and LQ subsets according to the image qualities, and learn them with classification loss and contrastive loss respectively. Then, we propose a proxy-updated real-time contrastive queue to support the face identification among LQ images.

A. Quality Partitioning

Considering the domain gap between the HQ and the LQ images, we partition the mixed-quality images into LQ and HQ subsets during training. In [15], [19], it is found that the feature norm is highly correlated with the image quality. In our approach, we use the feature norm as the quality indicator for data partitioning.

Let z_i denote the norm of the feature h_i , and μ_z , σ_z denote the mean and the standard deviation of z respectively. With these statistics, we could approach the normalized distribution of z .

$$\hat{z}_i = \left[\frac{z_i - \mu_z}{\sigma_z/c} \right]_{-1}^1, \quad (1)$$

where c is a scale parameter to keep most values of \hat{z} fall in $(-1,1)$. The gradient flow of \hat{z} is stopped during backpropagation to keep the norm from being optimized directly.

We partition the input mixed-quality data into HQ and LQ subsets based on the quality indicator \hat{z} using a quality threshold b . Classification loss and contrastive learning are then applied for the HQ and the LQ feature learning respectively. We define dependent but different quality partitioning strategies for these two methods according to their inherent properties. For classification loss L_{class} , each HQ feature is processed individually. We apply the quality comparison on the features in the same way, i.e. the features with quality indicator \hat{z}_i higher than b are fed into the classification loss. Contrastive loss L_{contra} focuses on the relationship between sample pairs. We take the representations of the augmented image and the original one as a positive contrastive pair (q, k) for learning. We define the quality indicator of the p -th positive pair as the lower one of its two components. The pairs whose quality indicators are lower than the quality threshold b are fed

into the contrastive learning. Specifically,

$$\hat{z}_p = \min(\hat{z}_p^q, \hat{z}_p^k), \quad (2)$$

$$\mathcal{L}_{classification} = \begin{cases} L_{class}, & \hat{z}_i > b \\ 0, & \text{else} \end{cases}, \quad (3)$$

$$\mathcal{L}_{contrastive} = \begin{cases} L_{contra}, & \hat{z}_p \leq b \\ 0, & \text{else} \end{cases}, \quad (4)$$

$$\mathcal{L} = \mathcal{L}_{classification} + \mathcal{L}_{contrastive}. \quad (5)$$

The outcomes of the two kinds of losses after quality partitioning are denoted as $\mathcal{L}_{classification}$ and $\mathcal{L}_{contrastive}$ respectively. We take their summation \mathcal{L} as the final loss.

B. Quality-guided joint training

The quality-invariant methods learn the LQ information from the corresponding features of the HQ images encoded by a pretrained model, and apply contrastive loss or the KL-divergence loss on all data. Different from them, we guide the model to learn the HQ and the LQ images simultaneously by applying classification loss and contrastive learning on different features according to their qualities. In this way, the supervision compromise between the HQ and the LQ features is minimized, and we can make full use of images with different qualities to train the encoder for the mixed-quality face recognition task. During the evaluation process, these two loss functions and their related structures will be dropped, and only a single feature extractor is reserved which can process images without distinction of qualities to enable mixed-quality face recognition.

1) *Classification loss for HQ features:* Margin-based softmax loss functions are widely used in face recognition and have achieved SOTA performances on the HQ evaluation datasets [1]–[3]. This kind of methods add a margin parameter to enlarge the distance between the matched image-identity pair and the unrelated ones to discriminate between different identities. Their loss could be shown as:

$$L = -\log \frac{e^{s \cdot f(\mathbf{h}_i, \mathbf{W}_{y_i})}}{e^{s \cdot f(\mathbf{h}_i, \mathbf{W}_{y_i})} + \sum_{j=1, j \neq y_i}^n e^{s \cdot f'(\mathbf{h}_i, \mathbf{W}_j)}}, \quad (6)$$

where $f(\mathbf{h}_i, \mathbf{W}_{y_i})$ and $f'(\mathbf{h}_i, \mathbf{W}_j)$ are two different functions to modulate the positive and the negative pair production of the feature $\mathbf{h} \in \mathbb{R}^d$, y_i denotes the corresponding

identity of \mathbf{h}_i , $\mathbf{W} \in \mathbb{R}^{d \times n}$ is the weight of the classifier with d being the feature dimension and n being the number of classes, and s is a scale parameter.

AdaFace [15] introduces two norm-related margin parameters g_{angle} and g_{add} to leverage the quality indicator with a predefined static margin m , where

$$g_{angle} = -m \cdot \hat{z}_i, \quad (7)$$

$$g_{add} = m \cdot \hat{z}_i + m, \quad (8)$$

$$f_{AdaFace}(\mathbf{h}_i, \mathbf{W}_{y_i}) = \cos(\theta_{y_i} + g_{angle}) - g_{add}, \quad (9)$$

$$f'_{AdaFace}(\mathbf{h}_i, \mathbf{W}_j) = \cos(\theta_j). \quad (10)$$

With further analysis on the gradient equations of L , a gradient scaling term (GST) g could be extracted to understand the effects of the cross-entropy softmax-based loss functions,

$$P_{y_i}^{(i)} = \frac{\exp(f(\cos \theta_{y_i}))}{\exp(f(\cos \theta_{y_i})) + \sum_{j=1, j \neq y_i}^n \exp(s \cos \theta_j)}, \quad (11)$$

$$g := (P_{y_i}^{(i)} - 1) \frac{\partial f(\cos \theta_{y_i})}{\partial \cos \theta_{y_i}}, \quad (12)$$

where $P_{y_i}^{(i)}$ denotes the probability of an input \mathbf{x}_i on its target class y_i . After substituting Equations 9 and 10 into Equation 6, we get the AdaFace loss which is taken as our classification loss L_{class} for the HQ feature learning.

2) *Contrastive learning for LQ features:* Although classification-based methods achieve SOTA performances on HQ face recognition, they are not suitable for learning LQ images. Fig. 2 visualizes the GST of AdaFace. It is obvious that LQ images gain more attentions during training. However, due to the low qualities, they lack the discriminative information for matching with thousands of abstractive identity features. To alleviate this problem, we apply contrastive learning on LQ features, which is based on the similarity between image instances and is more practical for learning the LQ data.

To comprehensively learn the LQ images and compose the positive contrastive pairs, we apply data augmentation on the HQ training images. Different from the popular contrastive methods which apply data augmentation on the images twice to compose the positive pairs, we only augment each image once to compose the pair with the original image. Also, we stop the gradients of HQ images

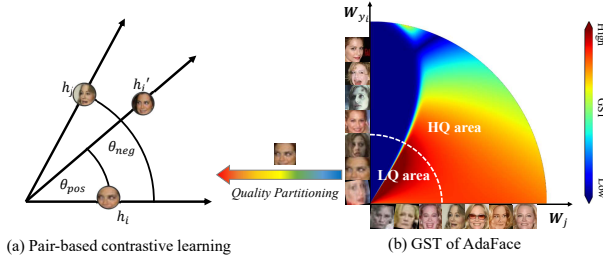


图 2. Illustration of contrastive learning for the LQ features. The distance between a data point and the origin denotes its feature norm. The GST shows that AdaFace puts a lot of attentions on the low-quality samples. The classification process matches the images with the abstract identity proxies, which is challenging when learning with the LQ data. We take the instance-level contrastive learning as shown in (a) on these images to relieve the learning burden.

during training to keep the qualities of their features. Let $(\mathbf{x}^q, \mathbf{x}^{k+})$ denote the augmented image and the original image respectively, and $(\mathbf{h}^q, \mathbf{h}^{k+})$ denote their features which compose a positive pair in contrastive learning.

Different augmentation methods are used including scaling (down-sampling to a lower resolution and then up-sampling to the original size), crop & random resize, rotation, and color distortion. AdaFace introduces hard data augmentations with low possibilities including cropping 20% of the original image, which make the augmented images unrecognizable. This might be acceptable when little emphasis is put on these data, but is not suitable in our situation where all LQ features are used for learning. We increase the possibility of data augmentation with lower-level distortion. In addition, we follow [6] to add JPEG compression artifacts to make the augmented images more realistic.

Finally, to make the above two learning paradigms better cooperate with each other, we employ the Supervised Contrastive Mask (SCM) introduced in CoReFace [32]. The pair-guided contrastive learning may push two features away as they are not generated from the same image, while the classification method tries to pull them together as they are from the same identity. To address this conflict, SCM removes all samples with the same identity of the given image from the negative feature pool. Let θ_{qk} , y_+ denote the angle between feature pairs and the label of the given image respectively, and Q , k_j denote the number of negative samples and the j^{th} feature among them. With a scale parameter s , the contrastive loss can be formulated

as

$$f_{contra}(\mathbf{h}^q, \mathbf{h}^k) = \cos(\theta_{qk}). \quad (13)$$

$$L_{contra} = -\log \frac{e^{s \cdot f_{contra}(\mathbf{h}^q, \mathbf{h}^{k+})}}{\sum_{j=1}^Q \mathbb{1}_{[y_j \neq y_+]} e^{s \cdot f_{contra}(\mathbf{h}^q, \mathbf{h}^{k_j})}}. \quad (14)$$

C. Proxy-Updated Real-Time Queue

Self-supervised contrastive learning usually needs a large comparison pool to serve as the negative candidates and push the training samples away from them. MoCo [17] defines the self-supervised contrastive learning as the comparison with a queue. A slowly progressing momentum encoder is proposed to produce keys and save them in a queue. The queries are produced by the training encoder. To make the positive and the negative keys comparable, all of them are produced by the momentum encoder. However, this operation prevents the training encoder from learning the LQ and the HQ image simultaneously as one of them is sent into a gradient-free encoder. Furthermore, as the momentum encoder is designed to be stable during large training intervals, this could lose track of the status of the encoder in joint training and produce whitened features.

To relieve the above problems, we propose a novel proxy-updated queue for contrastive learning, which is illustrated in Fig. 3. We take the genuine features encoded by the training model as the positive pairs. Since the training target of the cross-entropy softmax loss could constrain the corresponding feature-proxy pair to be similar, we update the negative features with their related proxies. Thus, we build a feature queue and a proxy queue. Specifically, we first encode all the original images and the augmented ones. Then, we update the feature queue and the proxy queue with the above features and their corresponding identity proxies. As the features in queue could be generated from several iterations ago, we update them to be similar to the features generated by the training encoder. They are added with the difference between the proxies in queue and the proxies in training. To provide a mixed-quality queue, we take both \mathbf{h}^q and \mathbf{h}^k to update it. With SCM, the comparisons between the identical features will not be introduced.

The cross-entropy softmax loss calculates the similarity between a feature and all of the identity proxies. To

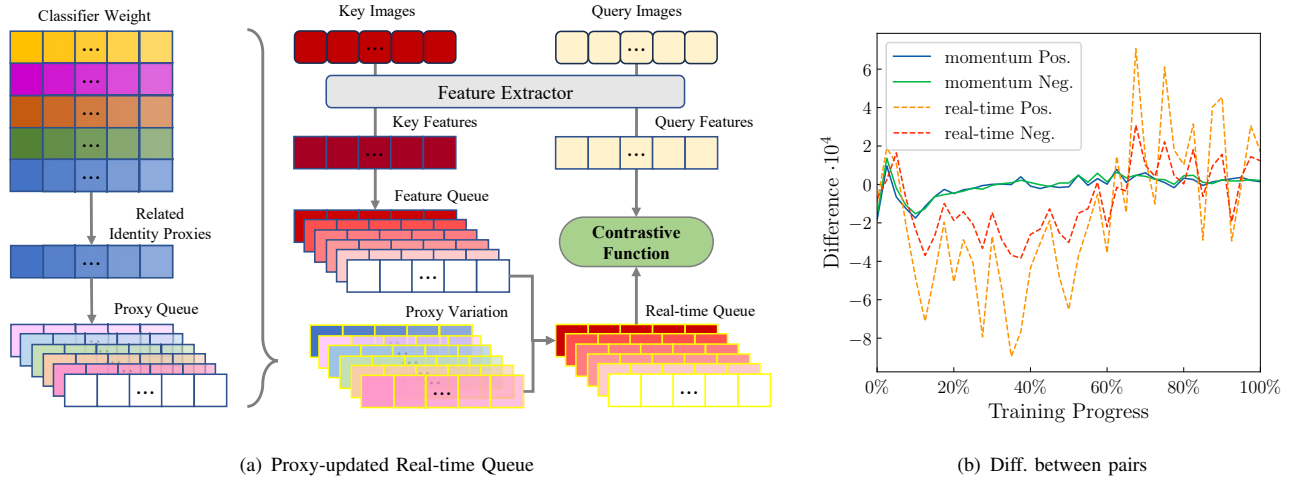


图 3. (a) The updating process of our proposed proxy-updated real-time queue; (b) The difference between the positive and the negative pairs with different queues. The momentum queue faces a strongly limited boundary with the steady and whitened features. With our queue, contrastive learning can take the features from the training encoder to compose the positive pairs.

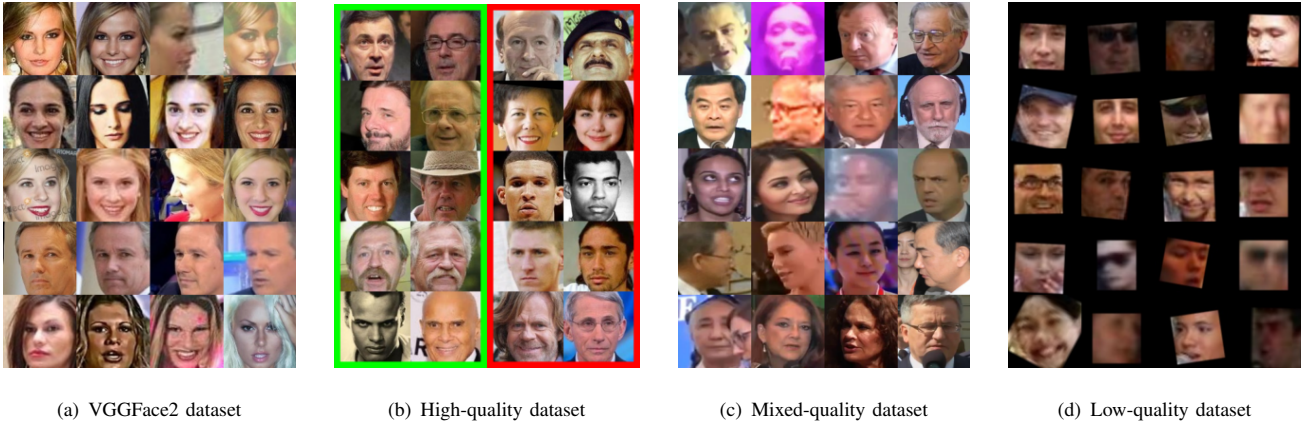


图 4. (a) Examples from the training dataset VGGFace2. Each row represents an identity. Identities are randomly selected from the dataset. (b) Examples from the high-quality dataset CALFW. The positive and the negative pairs are bounded with green and red. (c) Examples from the mixed-quality dataset IJB-B. (d) Examples from the low-quality dataset Tinyface. These face crops are generated from real images and very challenging.

provide a comparative magnitude of the pair relationship supervision, we set the size of the queue Q as the number of the identities n . Thus, the number of pair comparison used for contrastive learning keeps large in training. This enables the contrastive learning to discriminate among more images of mixed qualities for more effective learning.

IV. EXPERIMENTS

A. Datasets and Implementation Details

Datasets. We take VGGFace2 [33] as the training dataset. We evaluate our approach on three categories of datasets for different tasks. The first category is the HQ datasets including the widely-used LFW [34], AgeDB [35], CFP-FP [36], CPLFW [37], and CALFW [38]. The

second category is the LQ datasets including SCface dataset [12] and Tinyface dataset [23] which are usually used for the cross-resolution face recognition and the LQ face recognition respectively. SCface contains 130 identities captured by five surveillance cameras from 3 distances: 4.2m (d1), 2.6m (d2), and 1.0m (d3) for each identity. At each distance, the LQ images need to be matched with the HQ gallery images of the corresponding identities. Tinyface contains 5K identities with 153.5K distracting images. The images in the gallery and the probe sets are of low resolution. This setting and the natural LQ property make it a challenging dataset. For the third category, we choose IJB-B [39] as the mixed-quality dataset which contains about 1.8K identities with a total of 21.8K images and

表 1

ABLATIONS OF OUR PROPOSED FRAMEWORK. “HQ AVG.” DENOTES THE AVERAGE PERFORMANCE ON THE FIVE HQ DATASETS LFW, AGE DB, CFP-FP, CALFW, AND CPLFW. RANK-1 ACCURACY IS REPORTED FOR SCFACE AND TINYFACE. $TAR@FAR=10^{-4}$ IS REPORTED FOR IJB-B DATASET.

Setting	Augmentation	Quality Partition	Contrastive Learning	SCface			Tinyface	IJB-B	HQ Avg.
				d1	d2	d3			
Baseline	✗	✗	✗	61.69	97.23	99.85	63.98	90.26	95.62
A	AdaFace	✗	✗	66.77	99.08	100.0	67.49	74.07	95.76
B	AdaFace	✓	✗	68.15	98.46	100.0	67.03	70.45	95.73
C	AdaFace	✓	✓	70.92	98.62	99.69	66.58	90.59	95.71
D	Ours	✗	✗	88.62	99.23	100.0	68.80	89.50	95.46
E	Ours	✓	✓	92.31	99.54	100.0	69.85	91.05	95.43

55K unconstrained video frames. In the 1:1 verification task, there are about 10K positive matches and 8M negative matches. Examples of these datasets are shown in Fig. 4.

Implementation Details We follow the settings commonly used in recent works [15], [18], [19], [40], [41] to crop and resize the face images to 112×112 with five landmarks [1]. ResNet34 [42] is employed as the backbone model. We take AdaFace as the classification loss. Our framework is implemented in Pytorch [43]. We train the models on 4 NVIDIA A100 GPUs with a batch size of 512. All models are trained using SGD algorithm with an initial learning rate of 0.2. We set the momentum to 0.9 and the weight decay to 5×10^{-4} . The learning rate is divided by 10 after the 6th and the 9th epochs, and the training stops after 12 epochs. The margin parameter m is set to 0.4 as in AdaFace. We set the scale parameter s to 64 for both the classification loss and the contrastive loss.

B. Ablation on different training strategies

We investigate the effectiveness of the data augmentation strategy, the quality partitioning strategy (QP), and the contrastive learning (CL) with different categories of datasets in Table I. We demonstrate the effectiveness of our method on the LQ and the HQ datasets, and its generalizability to the mixed-quality dataset by jointly learning the LQ and the HQ data simultaneously with the guidance of image quality.

Baseline model: We use the plain ResNet [42] model with AdaFace [15] loss function as our baseline. No augmentation is used except for the random horizontal flip.

Setting A: To show the influence of the LQ data, we apply the data augmentations recommended in AdaFace,



图 5. Examples of the image pairs. Every two columns show several pairs of original images and the augmented images. The pairs in red box contain LQ images while the pairs in green box are all HQ images. Our quality partitioning strategy is capable of distinguishing low-quality images which are blurred or with obstacles, or contain only a part of face.

i.e. down-sampling with a following up-sampling, cropping with a small part, and random color distortion. All these three augmentations are applied with independent possibilities of 0.2.

Setting B: To show the negative influence of some extreme LQ data on the classification-based method, we apply quality partitioning on AdaFace loss by simply ignoring them (the contrastive learning is not presented).

Setting C: To show the effects of our proposed framework, we apply our QGFace with the AdaFace data augmentation. Both quality partitioning and contrastive learning are applied.

Setting D: To show the impacts of data augmentation, we apply our augmentation with the original AdaFace.

Setting E: Finally, we experiment our proposed method with quality partitioning, contrastive learning, and our data augmentation.

Table I shows the results of the above settings. The Baseline setting gains solid performances on the HQ datasets and IJB-B dataset. However, its performances on d1 of SCface dataset and Tinyface dataset are significantly poor. In setting A, when data augmentation empowers the original feature extractor to gain a higher performance on the HQ and the LQ datasets, the accuracy dramatically drops on the IJB-B dataset. This could be interpreted as a sacrifice of the discriminability. With a single classification supervision signal, the features become less distinctive to compromise between the HQ and the LQ images. On IJB-B dataset where the mixed-quality matching between images requires stronger discriminability, this feature degradation misleads the comparisons among a large number of faces with various qualities. Setting B shows that only dropping some LQ data helps the model gain better understanding on the LQ SCface dataset. This illustrates that some extremely low-quality data could mislead the training of the classification method. Even though AdaFace adapts the data importance according to the quality, it is harmful to learn these LQ data with classification loss. In Setting C, our approach relieves the degradation caused by data augmentation on IJB-B and outperforms the Baseline settings. Furthermore, the performance on SCface dataset is also improved. Meanwhile, our performance on the HQ data is only slightly dropped by less than 0.05% compared with Setting B. This shows the effectiveness of contrastive learning on the LQ data and the ability of our method to keep the comprehension on the HQ data. Setting D improves the performances on the LQ datasets and demonstrates the effectiveness of our data augmentation. Meanwhile, it is inferior to Setting C on IJB-B dataset which illustrates the effectiveness of our joint training framework again.

Finally, we implement Setting E, which is our proposed approach. AdaFace learns the hard-to-learn LQ images by classification loss with relatively less importance. Instead, we learn from them with contrastive learning. In this way, our model keeps learning the LQ images and gains more instance-level understanding from these less discriminative images. As a result, our method achieves the best results on almost all the evaluation datasets of different qualities. To ensure that most of the augmented




图 6. Similarity difference between QGFace and the Baseline on SCface. The histograms illustrate the similarity between probes and their related gallery images (matched pairs), where our method shows obvious advantages. The title of each sub-figure indicates the sub-dataset and the improvement of the difference on the mean similarity between the matched pairs and the most similar unmatched pairs.

LQ images are discriminable and realistic, we down-sample each image and then up-sample it to the original resolution. The discarding ratio, the degree of the brightness variations, and the contrast and saturation in color distortion are set to lower levels. In addition, we follow C. Kuo *et al.* [6] to add JPEG compression artifacts to make the image more realistic. With the new data augmentation, the augmented images become more understandable and compliant with the real-world low-resolution images. Fig. 5 shows some examples. The image pairs in red box are marked as LQ. They are harder to be identified than others in the training data (in green box) due to the blur, the obstacles, or the non-frontal positions in the images. Our quality partitioning strategy sends these LQ data from the original and the augmented images to contrastive learning, and this setting outperforms all the other settings on SCface and Tinyface. Furthermore, it also shows promising results on the IJB-B dataset which contains a huge number of image pairs of mixed qualities.

C. Ablation on the contrastive queue

Fig. 3 shows the average difference of the feature pairs with the momentum queue and our proxy-updated real-time queue. As the momentum queue is constructed with the features from a slowly-progressing momentum encoder, the features are steady and whitened. Both the positive and the negative pairs are limited in a tight boundary. While in real scenarios where the difference is computed with the features from the same training encoder, it shows a much larger gap between the query and the corresponding

表 II

ABLATION OF DIFFERENT QUEUE SETTINGS. “No” INDICATES APPLYING DATA AUGMENTATION AND QUALITY PARTITIONING WITHOUT CONTRASTIVE LOSS. “BATCH” INDICATES APPLYING OUR QUALITY-GUIDED TRAINING STRATEGY WITH THE SAMPLES IN BATCH AS THE NEGATIVE POOL FOR CONTRASTIVE LEARNING.

Queue	SCface			Tinyface	IJB-B	HQ Avg.	LFW	CFP-FP	AgeDB	CPLFW	CALFW
	d1	d2	d3								
No	87.85	99.38	100.0	68.24	90.57	95.47	99.45	96.00	94.20	94.67	93.05
Batch	87.54	99.38	100.0	68.67	87.38	95.28	99.55	95.96	94.38	93.85	92.68
Momentum	63.08	97.85	99.85	62.34	88.70	95.18	99.55	96.03	93.98	94.65	91.68
Proxy-updated	92.31	99.54	100.0	69.85	91.05	95.43	99.50	96.03	94.00	94.48	93.12

key. Our proxy-updated queue takes such features as the positive pairs in contrastive learning directly and updates the features in queue with the identity proxies. According to the large size of the queue compared with the training batch, the difference between the negative pairs shows smaller magnitudes compared with the positive features in our queue.

Table II shows the performances of the models with different queue settings. In the “No” situation, contrastive learning is not applied, while both the original and the augmented images are used for learning. In the “Batch” setting, the queue size is limited to the current batch. The number of feature-feature pairs in contrastive learning is extremely smaller than the feature-identity pairs in the classification loss. Thus, it stumbles on the IJB-B dataset which contains a large number of mixed-quality image pairs and needs higher discriminability. The momentum queue utilizes the whitened features to compose the positive and the negative pairs in contrastive learning. This misleads the optimization process in our joint training scenario. As a result, the performances of all datasets are lower than the setting without a queue. Our proxy-updated queue provides genuine positive pairs and a dynamic negative queue to keep the same magnitude on the number of pair similarity computation as in the classification process. This enables the contrastive loss to effectively learn from the LQ images and keep high discriminability. As a result, it outperforms the baseline and the other settings equipped with other queues on all benchmarks except a slight performance drop (0.04%) on the HQ datasets.

D. Ablation on the quality partitioning threshold

Table III shows the results of different quality partitioning thresholds. When no threshold is set, the classification loss and the contrastive loss supervise all the data at the same time. These two supervision signals guide the model to compromise between the features of the LQ and the HQ images, and the performances on SCface and HQ datasets are limited. To alleviate the trivial parameter selection, we scale \hat{z} to $(0, 1)$ and thus the value of the threshold b lies in the same range. With a lower threshold, more LQ data is learned with the classification loss, and the performances on SCface and Tinyface drop. A larger threshold gives up the utilization of the labels of the relatively-high-quality images, which then results in degraded features. Finally, we choose $b = 0.2$ in our implementation to balance the performances on the LQ and the HQ face images.

E. Analysis on the SCface dataset

Besides the quantitative analysis above, we further illustrate the difference between our QGFace and the Baseline on SCface in Fig.6. It is obvious that the matched pairs encoded by our QGFace are more similar than the Baseline method, which demonstrates the effectiveness of the quality partitioning strategy, the contrastive learning, and the selected data augmentation. As the evaluation process aims at finding an image among all of the candidates in gallery, the most similar unmatched image is also important for the successful match. We then calculate the statistics of the similarity between a probe and its matched gallery image, and the similarity between the probe and its most similar unmatched gallery image. The difference between the means of the two kinds of similarity scores is calculated, and the improvement of our method is noted

表 III

ABLATION OF DIFFERENT QUALITY PARTITIONING THRESHOLDS. “No” INDICATES THAT THE CLASSIFICATION LOSS AND THE CONTRASTIVE LOSS ARE SIMPLY ADDED TOGETHER TO IMPLEMENT JOINT TRAINING.

Threshold	SCface			Tinyface	IJB-B	HQ Avg.	LFW	CFP-FP	AgeDB	CPLFW	CALFW
	d1	d2	d3								
No	90.77	99.69	100.0	69.80	91.26	95.08	99.58	95.79	93.95	93.42	92.67
0.1	86.00	99.54	100.0	68.13	90.87	95.41	99.53	95.87	94.33	94.47	92.83
0.2	92.31	99.54	100.0	69.85	91.05	95.43	99.50	96.03	94.00	94.48	93.12
0.3	88.92	99.54	100.0	68.51	89.33	94.95	99.55	95.03	94.27	93.77	92.15

in the title of the sub-figures. When the distances become longer and the qualities of the images become lower, the improvement of QGFace gets larger. In the case of 4.2m, the improvement is even more than 600%, which clearly explains how our method improves the accuracy by more than 30% on SCface-d1 in Table I.

F. Comparison with SOTA methods

On SCface dataset. Table IV compares our method with the SOTA methods on SCface dataset. Many quality-invariant methods need a fine-tuning process to learn the target domain information. This breaks the general end-to-end framework and limits the evaluation confidence as only 130 identities and no distracting faces are provided in SCface. Furthermore, all these methods show lower accuracy scores in d3 compared with the pure AdaFace trained with the HQ data. AdaFace-Aug, which is the same as Setting A in Table I, achieves 100.0% by learning the relatively-high-quality augmented images. This verifies our hypothesis that the quality-invariant methods lose discriminative information during training. Our approach achieves the best performances on all the three distances with clear margins compared with the other approaches. This demonstrates the effectiveness of our proposed quality-guided joint training strategy for mixed-quality face recognition by simultaneously learning both the HQ and the LQ images.

On Tinyface dataset. Table V compares our method with other SOTA methods on Tinyface dataset. For the pure classification-based CurricularFace, even incorporated with a larger backbone model and a larger training dataset, it does not show an outstanding result compared with other methods. AdaFace puts relatively less emphasis on the hard LQ images, and outperforms CurricularFace. However, it is still under-competitive with the methods that focus on LQ

表 IV

RANK-1 ACCURACIES (%) OF DIFFERENT METHODS ON SCFACE. “FT” INDICATES THAT THE METHOD IS FINE-TUNED WITH A PART OF THE SCFACE DATASET. “ADAFACE-AUG” IS THE SAME AS SETTING A IN TABLE I.

Approach	Method	d1	d2	d3
Quality Invariant	T-C (IVC20) [10]	45.10	85.90	96.10
	FAN (ACCV19) [44]	62.00	90.00	94.00
	RAN (ECCV20) [13]	70.50	96.00	98.00
	S. -C. Lai <i>et al.</i> (APSIPA ASC21) [45]	79.70	95.70	98.20
	FAN-FT (ACCV19) [44]	77.50	95.00	98.30
	RAN-FT (ECCV20) [13]	81.30	97.80	98.80
	DDL (ECCV20) [14]	86.80	98.30	98.30
	RIFR (TBIOM20) [11]	88.30	98.30	98.60
Synthesis	SR-DCR (ICCP20) [27]	74.10	93.70	97.2
	M. Ullah <i>et al.</i> (ICET21) [28]	60.81	90.45	98.04
	IDEA-Net (TIFS22) [31]	90.76	98.50	99.25
Encoder-Only	NPT-Loss (TPAMI22) [46]	85.69	99.08	99.08
	AdaFace (CVPR22) [15]	61.69	97.23	99.85
	AdaFace-Aug (CVPR22) [15]	66.77	99.08	100.0
	PENB-FT (AAAI23) [20]	91.8	99.0	99.3
	Ours	92.31	99.54	100.0

face recognition. Our method partitions the data according to the image qualities and learns the LQ images with contrastive learning instead of harder classification process. The approach in RIFR [11] slightly outperforms our method by employing a distillation approach. This needs an extra pre-trained HQ face encoder to transfer its knowledge to the trained encoder, while our method does not have this preliminary requirement and is orthogonal with it. As we focus on training a feature extractor that is capable of dealing with the mixed-quality images, it is easy to integrate QGFace with other methods to further improve the performance. Furthermore, our method achieves a better average performance on SCface and Tinyface compared with RIFR [11], which demonstrates the effectiveness of our QGFace in recognizing mixed-quality face images under different scenarios.

表 V

RANK-1 ACCURACIES (%) OF DIFFERENT METHODS ON TINYFACE. * MEANS THAT THE METHOD IS TRAINED WITH RESNET100 ON MS1MV2 DATASET.

Approach	Method	Rank-1
Quality Invariant	QualNet50-LM (CVPR21) [47]	35.54
	MobileFaceNet (ICPR21) [48]	48.70
	T-C (IVC20) [10]	58.60
	URL (CVPR20) [49]	63.89
	RIFR (TBIOM20) [11]	70.40
	MIND-Resnet-50-FT (SPL21) [50]	66.82
Synthesis	CSRI-FT (ACCV18) [23]	44.80
	SRDA (ICoICT21) [29]	34.15
	CFSM (ECCV22) [30]	63.01
	IDEA-Net (TIFS22) [31]	68.13
Encoder-Only	CurricularFace(CVPR20)* [41]	63.68
	PENB-FT (AAAI23) [20]	39
	AdaFace (CVPR22) [15]	63.98
	AdaFace (CVPR22)* [15]	68.21
	Ours	69.85

V. CONCLUSION

We have presented our approach of quality-guided joint training strategy for mixed-quality face recognition. Based on quality partitioning, the high-quality data is supervised with the classification loss while the low-quality data is supervised with the contrastive loss. In this way, we can learn both the HQ and the LQ images simultaneously without any other network structures to cope with the mixed-quality faces in real applications. The proxy-updated real-time contrastive queue avoids the feature whitening problem in the momentum encoder queue and improves the discriminability on the LQ data. Our approach is as elegant as a single feature extractor, and can be easily integrated with other approaches such as super-resolution and distillation to further improve the performance.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (No. 61375016) and ECNU Multifunctional Platform for Innovation (001) Center for High Performance Computing.

参考文献

- [1] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive angular margin loss for deep face recognition," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2019, pp. 4690–4699.
- [2] 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 *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. IEEE Computer Society, 2018, pp. 5265–5274.
- [3] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song, "Sphereface: Deep hypersphere embedding for face recognition," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. IEEE Computer Society, 2017, pp. 6738–6746.
- [4] A. Sepas-Moghaddam, M. A. Haque, P. L. Correia, K. Nasrollahi, T. B. Moeslund, and F. Pereira, "A double-deep spatio-angular learning framework for light field-based face recognition," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 12, pp. 4496–4512, 2020.
- [5] W. Hu, W. Yan, and H. Hu, "Dual face alignment learning network for nir-vis face recognition," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 32, no. 4, pp. 2411–2424, 2022.
- [6] C. B. Kuo, Y.-T. Tsai, H.-H. Shuai, Y.-R. Yeh, and C.-C. Huang, "Towards understanding cross resolution feature matching for surveillance face recognition," *Proceedings of the 30th ACM International Conference on Multimedia*, 2022.
- [7] J. Chen, J. Chen, Z. Wang, C. Liang, and C.-W. Lin, "Identity-aware face super-resolution for low-resolution face recognition," *IEEE Signal Processing Letters*, vol. 27, pp. 645–649, 2020.
- [8] S.-C. Lai, C. He, and K.-M. Lam, "Low-resolution face recognition based on identity-preserved face hallucination," *IEEE International Conference on Image Processing (ICIP)*, pp. 1173–1177, 2019.
- [9] J. Zha and H. Chao, "Tcn: Transferable coupled network for cross-resolution face recognition*," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3302–3306, 2019.
- [10] F. V. Massoli, G. Amato, and F. Falchi, "Cross-resolution learning for face recognition," *Image and Vision Computing*, vol. 99, p. 103927, 2020.
- [11] S. S. Khalid, M. Awais, Z. Feng, C.-H. Chan, A. Farooq, A. Akbari, and J. Kittler, "Resolution invariant face recognition using a distillation approach," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 2, pp. 410–420, 2020.
- [12] M. Grgic, K. Delac, and S. Grgic, "Scface – surveillance cameras face database," *Multimedia Tools and Applications*, vol. 51, pp. 863–879, 2011.
- [13] H. Fang, W. Deng, Y. Zhong, and J. Hu, "Generate to adapt: Resolution adaption network for surveillance face recognition," in *Proc. of ECCV*, 2020.
- [14] Y. Huang, P. Shen, Y. Tai, S. Li, X. Liu, J. Li, F. Huang, and R. Ji, "Improving face recognition from hard samples via distribution distillation loss," in *Proc. of ECCV*, 2020.
- [15] M. Kim, A. K. Jain, and X. Liu, "Adaface: Quality adaptive margin for face recognition," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. IEEE, 2022, pp. 18 729–18 738.
- [16] S.-M. Yang, W. Deng, M. Wang, J. Du, and J. Hu, "Orthogonality loss: Learning discriminative representations for face recognition," *IEEE Trans. Circuits Syst. Video Technol.*, 2021.
- [17] K. He, H. Fan, Y. Wu, S. Xie, and R. B. Girshick, "Momentum contrast for unsupervised visual representation learning," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2020, pp. 9726–9735.
- [18] Y. Kim, W. Park, and J. Shin, "BroadFace: Looking at tens of thousands of people at once for face recognition," in *Proc. of ECCV*, 2020.
- [19] Q. Meng, S. Zhao, Z. Huang, and F. Zhou, "Magface: A universal representation for face recognition and quality assessment," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2021, pp. 14 225–14 234.

- [20] X. Ling, Y. Lu, W. Xu, W. Deng, Y. Zhang, X. Cui, H. Shi, and D. Wen, "Dive into the resolution augmentations and metrics in low resolution face recognition: A plain yet effective new baseline," in *AAAI workshop*. AAAI Press, 2023.
- [21] 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 *Proc. of ECCV*, 2018.
- [22] L. Xu and Z. Gajic, "Improved network for face recognition based on feature super resolution method," *International Journal of Automation and Computing*, vol. 18, pp. 915 – 925, 2021.
- [23] Z. Cheng, X. Zhu, and S. Gong, "Low-resolution face recognition," in *Asian Conference on Computer Vision*, 2018.
- [24] X. Shan, Y. Lu, Q. Li, and Y. Wen, "Model-based transfer learning and sparse coding for partial face recognition," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 31, no. 11, pp. 4347–4356, 2021.
- [25] M. Jian and K.-M. Lam, "Simultaneous hallucination and recognition of low-resolution faces based on singular value decomposition," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 25, no. 11, pp. 1761–1772, 2015.
- [26] A. Bulat, J. Yang, and G. Tzimiropoulos, "To learn image super-resolution, use a gan to learn how to do image degradation first," in *Proc. of ECCV*, 2018.
- [27] L. Ren, W. Deng, L. Wang, C. Mao, Y. Jiang, H. Jia, and J. Li, "Low-resolution face recognition method combining super-resolution and improved dcr model," *Proceedings of the 2020 9th International Conference on Computing and Pattern Recognition*, 2020.
- [28] M. M. I. Ullah, A. Hamza, I. A. Taj, and M. Tahir, "Low resolution face recognition using enhanced srgan generated images," *International Conference on Emerging Technologies (ICET)*, pp. 1–6, 2021.
- [29] J. C. Tan, K. M. Lim, and C. P. Lee, "Enhanced alexnet with super-resolution for low-resolution face recognition," *International Conference on Information and Communication Technology (ICOICT)*, pp. 302–306, 2021.
- [30] F. Liu, M. Kim, A. K. Jain, and X. Liu, "Controllable and guided face synthesis for unconstrained face recognition," in *Proc. of ECCV*, vol. 13672. Springer, 2022, pp. 701–719.
- [31] C. Low and A. B. J. Teoh, "An implicit identity-extended data augmentation for low-resolution face representation learning," *IEEE Trans. Inf. Forensics Secur.*, vol. 17, pp. 3062–3076, 2022. [Online]. Available: <https://doi.org/10.1109/TIFS.2022.3201374>
- [32] Y. Song and F. Wang, "Coreface: Sample-guided contrastive regularization for deep face recognition," *CoRR*, vol. abs/2304.11668, 2023.
- [33] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "Vggface2: A dataset for recognising faces across pose and age," *IEEE International Conference on Automatic Face & Gesture Recognition, FG*, pp. 67–74, 2018.
- [34] G. B. Huang, M. A. Mattar, T. L. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments," in *Tech. Rep.*, 2007.
- [35] S. Moschoglou, A. Papaioannou, C. Sagonas, J. Deng, I. Kotsia, and S. Zafeiriou, "Agedb: The first manually collected, in-the-wild age database," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2017.
- [36] S. Sengupta, J. Chen, C. D. Castillo, V. M. Patel, R. Chellappa, and D. W. Jacobs, "Frontal to profile face verification in the wild," in *WACV*, 2016.
- [37] T. Zheng and W. Deng, "Cross-pose lfw : A database for studying cross-pose face recognition in unconstrained environments," in *Tech. Rep.*, 2018.
- [38] T. Zheng, W. Deng, and J. Hu, "Cross-age LFW: A database for studying cross-age face recognition in unconstrained environments," *CoRR*, 2017.
- [39] C. Whitelam, E. Taborsky, A. Blanton, B. Maze, J. C. Adams, T. Miller, N. D. Kalka, A. K. Jain, J. A. Duncan, K. Allen, J. Cheney, and P. Grother, "IARPA janus benchmark-b face dataset," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2017.
- [40] X. Wang, S. Zhang, S. Wang, T. Fu, H. Shi, and T. Mei, "Mis-classified vector guided softmax loss for face recognition," in *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 2020, pp. 12 241–12 248.
- [41] Y. Huang, Y. Wang, Y. Tai, X. Liu, P. Shen, S. Li, J. Li, and F. Huang, "Curricularface: Adaptive curriculum learning loss for deep face recognition," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. IEEE, 2020, pp. 5900–5909.
- [42] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*. IEEE Computer Society, 2016, pp. 770–778.
- [43] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," in *NeurIPS Workshop*, 2017.
- [44] X. Yin, Y. Tai, Y. Huang, and X. Liu, "Fan: Feature adaptation network for surveillance face recognition and normalization," in *Asian Conference on Computer Vision*, 2019.
- [45] S.-C. Lai and K.-M. Lam, "Deep siamese network for low-resolution face recognition," *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pp. 1444–1449, 2021.
- [46] S. S. Khalid, M. Awais, Z. Feng, C.-H. Chan, A. Farooq, A. Akbari, and J. Kittler, "Npt-loss: Demystifying face recognition losses with nearest proxies triplet," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [47] I. Kim, S. Han, J. Baek, S. Park, J. Han, and J. Shin, "Quality-agnostic image recognition via invertible decoder," in *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2021, pp. 12 257–12 266.
- [48] Y. Martínez-Díaz, H. M. Vazquez, L. S. Luevano, L. Chang, and M. González-Mendoza, "Lightweight low-resolution face recognition for surveillance applications," *International Conference on Pattern Recognition (ICPR)*, pp. 5421–5428, 2021.
- [49] Y. Shi, X. Yu, K. Sohn, M. Chandraker, and A. K. Jain, "Towards universal representation learning for deep face recognition," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR*. IEEE, 2020, pp. 6816–6825.
- [50] C.-Y. Low, A. Teoh, and J. Park, "Mind-net: A deep mutual information distillation network for realistic low-resolution face recognition," *IEEE Signal Processing Letters*, vol. 28, pp. 354–358, 2021.



Youzhe Song received the B.E. degree in School of Computer Science and Technology from the DongHua University (DHU), Shanghai, China, in 2021. He is currently pursuing the M.E. degree at the School of Computer Science and Technology from the East China Normal University (ECNU), Shanghai, China. His research interests include face recognition and open-set retrieval.



Feng Wang received his PhD in Computer Science from the Hong Kong University of Science and Technology in 2007 and BsC from Fudan University, China, in 2001 respectively. Before joining East China Normal University as an associate professor in the Dept. of Computer Science and Technology, he was a research fellow in City University of Hong Kong and Institute Eurecom, France. His research interests include multi-media information retrieval, pattern recognition and IT in education.