

ARoFace: Alignment Robustness to Improve Low-Quality Face Recognition

Mohammad Saeed Ebrahimi Saadabadi*, Sahar Rahimi Malakshan*,
Ali Dabouei, and Nasser M. Nasrabadi

West Virginia University

*Equal Contribution

{me00018, sr00033, Ad0046}@mix.wvu.edu, nasser.nasrabadi@mail.wvu.edu

Abstract. Aiming to enhance Face Recognition (FR) on Low-Quality (LQ) inputs, recent studies suggest incorporating synthetic LQ samples into training. Although promising, the quality factors that are considered in these works are general rather than FR-specific, *e.g.*, atmospheric turbulence, resolution, *etc.* Motivated by the observation of the vulnerability of current FR models to even small Face Alignment Errors (FAE) in LQ images, we present a simple yet effective method that considers FAE as another quality factor that is tailored to FR. We seek to improve LQ FR by enhancing FR models' robustness to FAE. To this aim, we formalize the problem as a combination of differentiable spatial transformations and adversarial data augmentation in FR. We perturb the alignment of the training samples using a controllable spatial transformation and enrich the training with samples expressing FAE. We demonstrate the benefits of the proposed method by conducting evaluations on IJB-B, IJB-C, IJB-S (+4.3% Rank1), and TinyFace (+2.63%). <https://github.com/msed-Ebrahimi/ARoFace>

Keywords: Unconstrained Face Recognition · Face Alignment · Low-Quality Input · Adversarial Data Augmentation

1 Introduction

Excellence in Face Recognition (FR) is attributed to large-scale training datasets [21, 92], advanced deep networks [22, 24, 57], and angular criteria [11, 36, 42, 62]. However, current FR models experience performance failure in practical scenarios where images are acquired from long-range distances, *i.e.*, Low-Quality (LQ) or unconstrained images [2, 3, 5, 31, 54]. Specifically, the reported performance [41, 55] of the state-of-the-art (SOTA) models on IJB-S [31] and TinyFace [5], *i.e.*, LQ benchmarks, are about 30% lower than LFW [25], *i.e.*, High-Quality (HQ) benchmark. This deficiency stems from the distribution gap between the training and LQ testing data, *i.e.*, lack of a sufficient number of LQ instances during training [18, 41, 54, 58, 65, 84]. An intuitive solution would be to construct a large-scale dataset with an adequate amount of LQ samples which is impractical due to data acquisition costs and privacy concerns [41, 60].

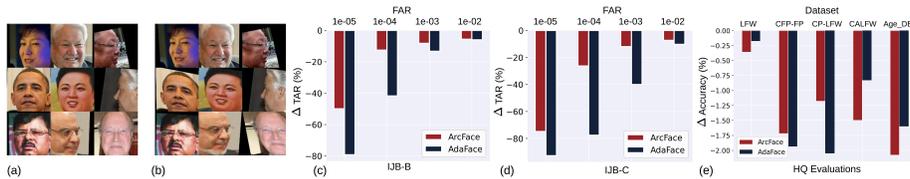


Fig. 1: Visual comparison of aligned (a) and alignment-perturbed (b) samples from the IJB-B dataset. (c, d, e) The performance difference between aligned inputs and those with slight FAE. Models exhibit robustness to FAE in HQ samples but suffer significant performance drops in LQ faces, with over 50% reduction in TAR@FAR = $1e-5$. Results from two distinct ResNet-100 trained on MS1MV3 using ArcFace/AdaFace objective.

Considering the components of a practical FR system, *i.e.*, face detection, face alignment, and face recognition model [70, 86], available LQ FR methods primarily have focused on how LQ inputs affect the recognition model [4, 18, 23, 40, 41, 56, 58, 65, 68, 78]. These studies can be categorized into Facial Image Enhancement (FIE) [4, 20, 23, 40, 43, 68, 78, 83], and Common Space Mapping (CSM) [18, 41, 56, 58, 65, 84]. FIE estimates HQ face from LQ counterpart and then performs FR, yet faces the challenge of ill-posedness, *i.e.*, multiple HQ outcomes exist for a single LQ face [26, 56, 75, 80]. CSM studies [18, 56, 58, 65, 84] seek to map LQ and HQ instances to a common embedding by employing image degradations as data augmentation during training. However, recent CSM works lack practicality due to deficiency in expressing real-world LQ [3, 47, 47, 48], entailing high computational resources [50], identity preservation challenge [1, 39, 56, 66, 76], and constrained of *a priori* defined target data [48, 87].

We extend our focus beyond the face recognition model to explore the impact of LQ inputs on face detection and alignment. It has been shown that face detector accuracy diminishes with LQ inputs, leading to Face Alignment Error (FAE) [10, 35, 59, 70, 85]. We investigate the impact of FAE on FR performance by manually adding small random spatial perturbations including scaling, rotation, and translation, to the original aligned evaluation faces, as shown in Figures 1a, and b. The results in As depicted in Figures 1c, d, and e, the performance gap of a given model between alignment-perturbed and aligned samples. Based on these observations, we found that while SOTA FR networks are robust against FAE in HQ faces, they are susceptible to FAE in LQ samples, *e.g.*, more than 50% drop in True Acceptance Rate (TAR) at a False Acceptance Rate (FAR) of $1e-5$ (TAR@FAR= $1e-5$) on IJB-B and IJB-C.

Therefore, we contend that FAE plays a crucial role in the failure of FR models on LQ faces. Current CSM methods [18, 41, 56, 84] enrich training with synthetic LQ faces yet ignore the FAE as another degradation component. Consequently, the current FR models are extremely susceptible to the intersection of FAE and LQ input. A potential remedy is end-to-end training of both recognition and detection networks. However, this is limited by high computational demands and the lack of large datasets with both types of labels. A few studies tried to establish alignment-free FR by incorporating an alignment module into the recognition model [28, 70, 88, 89]. However, the estimated spatial transforma-

tion expresses coarse geometric information as a holistic parametric model and falls short of improving practical LQ FR [35, 41, 54].

In an orthogonal direction, we consider FAE as another image degradation factor and aim to develop an FR model robust against FAE, dubbed Alignment Robust Face (ARoFace) recognition. Our proposal is inspired by the adversarial data augmentation that enhances model robustness through integrating adversarial components into the training process [19, 72]. Concretely, ARoFace employs a differentiable spatial transformation [28] to adversarially perturb the alignment of training samples and enrich the training with samples expressing FAE.

We craft samples expressing FAE by applying spatial transformations including scaling, rotation, and translation, to the benign (aligned) sample. We employ a global transformation grid [28] that avoids face deformation to compute faithful and beneficial instances with meaningful spatial variations. Moreover, a randomized adversarial budget is used to improve the diversity of the crafted samples, helping to attain a variety of spatial transformations for a single input.

Compared to recent face degradation methods [41, 54] our approach imposes negligible trainable parameters to the training pipeline, *i.e.*, a maximum of nine additional parameters. Moreover, the proposed method avoids the class distortion problem, *i.e.*, changing the identity in the output face, of GAN-based face degradation by not altering the intensity of the faces and merely modifying the position of pixels [1, 39, 56, 66, 76]. Thus, the resulting misaligned face completely lies in the manifold of faithful faces for the FR training [19, 48, 91]. Furthermore, ARoFace circumvents the two-stage optimization and the necessity of accessing target data in previous GAN-based studies [41, 54]. It is worth noting that our method does not require any face or landmark detection module, and can be readily integrated into the training pipeline of an arbitrary FR model to improve the generalization across LQ evaluations. In summary, the contributions of the paper are as follows:

- We introduce FAE as an image degradation factor tailored for FR which has previously been ignored in LQ FR studies.
- We propose an optimization method that is specifically tailored to increase the FR model robustness against FAE.
- We show that the proposed optimization can greatly increase the FR performance in real-world LQ evaluations such as IJB-S and TinyFace. Moreover, our framework achieves these improvements without sacrificing the performance on datasets with both HQ and LQ samples such as IJB-B and IJB-C.
- We empirically show that the proposed method is a plug-and-play module, providing an orthogonal improvement to SOTA FR methods.

2 Related Works

2.1 Low-quality Face Recognition

Approaches to tackle LQ FR can be categorized into: 1) Facial Image Enhancement (FIE) [4, 30, 43, 83], and 2) Common Space Mapping (CSM) [18, 58, 65, 84].

FIE methods try to solve the inverse problem of retrieving HQ samples from their LQ counterparts and then perform FR [26, 75, 80]. Despite the success in enhancing visual quality metrics such as PSNR [67], SSIM [67], and LPIPS [29], they fail to increase the FR performance on real-world testing benchmarks since estimating HQ face from an LQ counterpart is an ill-posed problem, *I.e.*, multiple HQ faces exist for single LQ input [34, 41, 79]. Conventional CSM methods try to find joint embedding for (HQ, LQ) pairs using sample-wise supervision. However, these approaches suffer from convergence issues when integrated into large-scale FR training framework [64].

Recently, Shi *et al.* proposed to employ an unlabeled LQ face dataset to establish a GAN as an image degradation module. In the follow-up work [41], Liu *et al.* incorporate feedback from the FR network to the degradation so it can control the amount of image degradation to be applied to each samples. However, the requirement of *a priori* selected target dataset limits their practical usability [48]. Moreover, GANs are cumbersome to deploy during large-scale training and can introduce class distortion issues [56, 66, 76]. Also, developing a proper generative model heavily relies on prior knowledge [1, 39, 65, 76].

2.2 Adversarial Data Augmentation in Face Recognition

Employing adversarial images as additional training data has been extensively surfed in various deep learning applications [19, 38, 44, 72, 73]. Broadly, adversarial data augmentation in FR intends to enhance the model’s resilience against certain adversarial components in the inputs [7, 8, 37, 52, 82]. Sharif *et al.* [52] devised real-world adversarial samples by adding printed glasses to the face images, while Komkov *et al.* [37] explored adversarial hats. Deb *et al.* [8] initially leveraged GANs for synthesizing adversarial samples. Subsequently, Yin *et al.* [82] introduced an adversarial makeup generation framework, while Dong *et al.* [15, 16] harnessed generative models for creating adversarial attributes. Dabouei *et al.* employed spatial transformation on facial landmarks to produce geometrically-perturbed adversarial faces. Liu *et al.* [41] employed the domain-translation property of GAN combined with adversarial training to boost the LQ FR. However, the necessity of data collection and pre-defined target data make these methods extremely specific to a particular scenario and reduce the practicality.

2.3 Face Alignment and Learnable Spatial Transformation

The pioneering work of Jaderberg *et al.* [28] introduced a differentiable layer that performs Spatial Transformation (ST) to an input image or feature map. Owing to its intuitive transparency, the ST layer has inspired a multitude of subsequent studies in applications including, but not limited to, classification, dense correspondence matching, and FR [33, 74, 90].

Face alignment focuses on matching faces to a unified template, thereby minimizing variations in geometry that are unrelated to identity. This typically involves a 2D affine transformation to map facial landmarks to a specific 2D

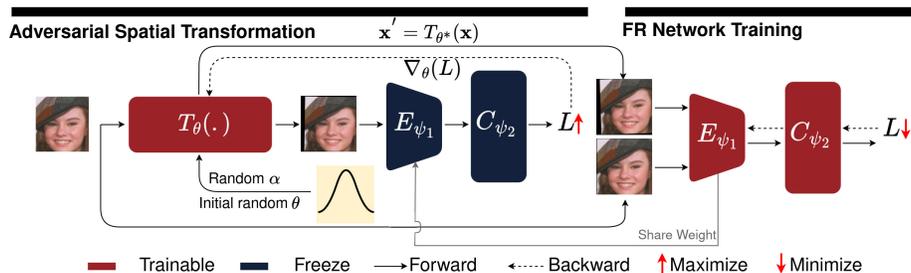


Fig. 2: Overview of proposed method. Each training iteration is composed of two steps. The adversarial spatial transformation finds $\theta^* = (\varphi^*, \Delta u^*, \Delta v^*, \lambda^*)$ for each instance in the batch based on the feedback from the FR network to produce hard but faithful samples, *i.e.*, maximization of L . Then the FR network is trained using a batch of adversarial and original samples, *i.e.*, minimization of L .

template [12, 42, 61, 63]. Efforts to create an end-to-end FR system have led to investigations into learning flexible, non-rigid transformations for establishing an alignment-free FR system, *i.e.*, integrating alignment into recognition module [70, 89]. The approach by Wu *et al.* [70] involved a recursive spatial transformer for complex transformation learning. Similarly, Zhou *et al.* [89] employed locally estimated homography transformations through a rectification network for face correction. These techniques aim to achieve alignment independence by concurrently learning alignment with the recognition process in a seamless, end-to-end manner. Taking a different approach, Xu *et al.* [74] explored optimizing the target template for alignment, rather than adjusting alignment parameters. However, their practical use is often hindered by computational demands and a potential decrease in the discriminatory capacity of the recognition network [35].

3 Proposed Method

3.1 Notation

In this paper, we use lowercase letters (*e.g.*, x) to denote scalars, lowercase boldface (*e.g.*, \mathbf{x}) to denote vectors, uppercase letters (*e.g.*, X) to denote functions, and uppercase calligraphic symbols (*e.g.*, \mathcal{X}) to denote sets.

Let $\mathcal{D} = \{(\mathbf{x}_n, y_n)\}_{n=1}^{n_0}$ be the training dataset, consisting of n_0 aligned faces from c classes. Furthermore, $F_{\psi} = C \circ E$ denotes a deep neural network with trainable parameter $\psi = [\psi_1, \psi_2]$ where $E_{\psi_1}(\cdot): \mathbb{R}^a \rightarrow \mathbb{R}^d$ is the deep feature extractor that maps an input face $\mathbf{x} \in \mathbb{R}^a$, to a d -dimensional representation $\mathbf{z} \in \mathbb{R}^d$, and $C_{\psi_2}: \mathbb{R}^d \rightarrow \mathbb{R}^c$ is a parametric classifier, *i.e.*, conventionally hyperspherical classifier in FR [11, 42, 62], with parameter ψ_2 that maps \mathbf{z} to a probability distribution over c classes. For brevity and convenience of presentation, all representations are ℓ_2 -normalized.

3.2 Preliminaries

We propose to enhance LQ FR by increasing robustness to Face Alignment Error (FAE) since FAE in LQ images is inevitable [10, 35, 59, 70, 85] and current FR

models are severely susceptible to FAE, as discussed in Section 1. To this end, the proposed method considers FAE as another image quality factor, *I.e.*, tailored to FR, and leverages adversarial data augmentation combined with differentiable spatial transformation to enrich the training with samples expressing FAE. In the next two sections, we introduce adversarial data augmentation and differentiable spatial transformation, respectively, as prerequisites to the proposed method. Then, we provide a detailed description of our method.

Adversarial Data Augmentation aims to improve training without requiring additional data [7, 44, 49]. Specifically, it seeks to craft hard training instances (adversarial samples) from the original sample \mathbf{x} (benign):

$$\theta^* = \arg \max_{\theta} L(F_{\psi}; T_{\theta}(\mathbf{x}), y), \quad (1)$$

where $T_{\theta}(\cdot)$ is an augmentation function with parameter θ , and L is the learning objective, *I.e.*, empirical risk. This optimization can be solved through k -steps Projected Gradient Descent (PGD) [44]:

$$\theta_{(k+1)} \leftarrow \text{proj}_{\mathcal{S}} \left(\theta_{(k)} + \alpha \text{sgn} \left(\nabla_{\theta_{(k)}} L(F_{\psi}; T_{\theta_{(k)}}(\mathbf{x}), y) \right) \right), \quad (2)$$

where $\text{sgn}(\cdot)$ denotes the sign function, α is the PGD step size, \mathcal{S} is set of allowed perturbation and $\text{proj}_{\mathcal{S}}(\mathbf{s})$ projects \mathbf{s} back into \mathcal{S} , $\text{proj}_{\mathcal{S}}(\mathbf{s}) = \arg \min_{\mathbf{s}' \in \mathcal{S}} \|\mathbf{s} - \mathbf{s}'\|_2$. Typically, \mathcal{S} is defined as ℓ_p -norm ball center at \mathbf{x} with radius ρ :

$$\mathcal{S} = \{\theta \mid \|T_{\theta}(\mathbf{x}) - \mathbf{x}\|_p \leq \rho\}. \quad (3)$$

Spatial Transformation. Here, we detail the spatial transformation applied to every channel of input data, $\mathbf{x} \in \mathbb{R}^q$, where q is $3 \times h \times w$ for RGB input. Let the row and column indices (i, j) be as points $(u, v) \in \mathbb{R}^2$, where the u -axis and the v -axis are horizontal and vertical axis, respectively. Furthermore, $P_u(j) = j - \frac{w-1}{2}$ and $P_v(i) = \frac{h-1}{2} - i$ convert zero-indexed (i, j) to u, v coordinates. T_{θ} is an invertible affine transformation with parameter $\theta = (\varphi, \Delta u, \Delta v, \lambda)$, where $\varphi \in [0, 2\pi]$ denotes the rotation angle, $\Delta u \in \mathbb{R}$, $\Delta v \in \mathbb{R}$ denote horizontal and vertical shifts, respectively, and $\lambda \in \mathbb{R}$ and $\lambda > -1$ denotes the scaling factor [77]. For each location $(P_u(j), P_v(i))$ the coordinate that maps to this location under T_{θ} can be obtained as:

$$(u', v') = T_{\theta}^{-1}(P_u(j), P_v(i)), \quad (4)$$

where T_{θ}^{-1} is the inverse transformation. (u', v') may not align with integer-valued pixel indices in the input image, hence, we can utilize the bi-linear interpolation kernel of Jaderberg *et al.* [28]:

$$I_x(u, v) = \sum_{i=0}^{h-1} \sum_{j=0}^{w-1} x_{i,j} \cdot \max(0, 1 - |v - P_v(i)|) \cdot \max(0, 1 - |u - P_u(j)|), \quad (5)$$

Algorithm 1 ARoFace

Input: Dataset \mathcal{D} , \bar{f} for conventional five landmarks in FR, number of PGD steps k , distribution Θ for θ , distribution N for α , and the total training iteration t

Output: FR network F_ψ

- 1: Initialize F_ψ
- 2: **for** $i = 0 \dots t - 1$ **do** ▷ outer optimization
- 3: Sample \mathcal{B} from \mathcal{D}
- 4: Sample $\theta \sim \Theta$
- 5: Sample $\alpha \sim N$
- 6: **for** $j = 1 \dots k$ **do** ▷ inner optimization for k steps
- 7: Compute $l = L(F_\psi; T_\theta(\mathcal{B}_x), \mathcal{B}_y)$
- 8: Update $\theta_{j+1} \leftarrow \theta_j + \alpha \text{sgn}(\nabla_{\theta_j}(l))$
- 9: $\theta^* \leftarrow \text{proj}_{\mathcal{S}}(\theta_{j+1})$
- 10: **end for**
- 11: Compute $l_1 = L(F_\psi; T_{\theta^*}(\mathcal{B}_x), \mathcal{B}_y)$ ▷ FR objective alignment-perturbed faces
- 12: Compute $l_2 = L(F_\psi; \mathcal{B}_x, \mathcal{B}_y)$ ▷ FR objective aligned faces
- 13: Update $\psi \leftarrow \psi - \eta \nabla_\psi(l_1 + l_2)$
- 14: **end for**

where $x_{i,j}$ denotes pixel value in the i -th row and j -th column of an arbitrary channel in \mathbf{x} . Thus the value of each pixel in the transformed image \mathbf{x}' is:

$$x'_{i,j} = I_x(T_\theta^{-1}(P_u(j), P_v(i))), \quad (6)$$

detailed derivation can be found in [28]. With a slight abuse of notation, we term $\mathbf{x}' = T_\theta(\mathbf{x})$ as the transformed version of \mathbf{x} under the transformation T_θ .

3.3 Alignment Robust Face Recognition

As observed in Figure 1, conventional FR networks are susceptible to even mild Face Alignment Errors (FAE) in LQ faces. This vulnerability is a primary cause for LQ FR failure since FAE is inevitable in LQ images [10, 35, 59, 70, 85]. Previous LQ FR works [4, 18, 41, 54, 65, 83, 84] have ignored FAE and solely tried to make the FR model robust against general image quality factors that are not tailored to FR, *e.g.*, resolution, atmospheric turbulence, blur, *etc.* We consider FAE as another degradation component and seek to improve LQ FR performance by making the FR model robust against FAE. To this end, the intuitive solutions are either to use unaligned faces or to apply random spatial transformations to the aligned samples. However, it has been shown that random data augmentation is ineffective in most FR evaluations [36, 41, 51].

We aim to make the FR network robust to FAE through enriching training with samples expressing FAE. To this aim, we leverage adversarial data augmentation (detailed in Section 3.2) combined with differentiable spatial transformation $T_\theta(\cdot)$ [28] (composition of scaling, rotation, and translation as detailed in Section 3.2) to craft adversarial samples expressing FAE [28]. Formally, we find the θ that crafts hard samples for the FR model, *i.e.*, maximizing the FR

objective function. Then, the FR network $F_\psi(\cdot)$ is trained using batch formed from benign and adversarial samples, *I.e.*, minimization of FR objective. Hence, our main optimization can be formulated as:

$$\arg \min_{\psi} \frac{1}{|\mathcal{D}|} \sum_{(\mathbf{x}, y) \in \mathcal{D}} \left[L(F_\psi; \mathbf{x}, y) + \arg \max_{\theta} L(F_\psi; T_\theta(\mathbf{x}), y) \right], \quad (7)$$

where the maximization is solved using PDG as explained in Equation 2 and L is an arbitrary FR objective, *I.e.*, ArcFace [11].

In crafting the adversarial sample $\mathbf{x}' = T_{\theta^*}(\mathbf{x})$, *I.e.*, maximization in Equation 7, we do not want \mathbf{x}' to be indistinguishable from \mathbf{x} , but it should lie on the manifold of valid training instances. However, the maximization may destroy all features of the input image. For instance, T_{θ^*} may zero out the image, making learning infeasible. Typically, the ℓ_p -norm of the difference between benign and adversarial samples is employed to define the allowed perturbation set \mathcal{S} , as shown in Equation 3. This convention is based on the fact that when two images are close in a ℓ_p -norm, they are visually similar [19]. However, its converse does not always hold, *I.e.*, two images that are distant in a ℓ_p -norm, can also be visually similar. For instance, large ℓ_2 -norm results from single pixel translation [17]. Thus ℓ_p -norm constrained \mathcal{S} is not applicable here.

Alternatively, we utilize the per-landmark flow (displacement) vector $\mathbf{f} = (u' - P_u(j), v' - P_v(i))$ [71], *I.e.*, the vector from a position of a landmark in the \mathbf{x}' to its corresponding position in the \mathbf{x} , to define the \mathcal{S} :

$$\mathcal{S} = \{\theta \mid \sum_{\mathbf{p} \in \mathcal{P}} \|\mathbf{f}_{\mathbf{p}}^\theta\|_2 \leq \sum_{\mathbf{p} \in \mathcal{P}} \bar{f}_{\mathbf{p}}\}, \quad (8)$$

where \mathcal{P} is the set of five landmarks typically employed in FR face alignment [7, 11, 36, 42, 62], and $\bar{f}_{\mathbf{p}}$ is the upper bound to the norm of flow vector corresponding to \mathbf{p} . Typically, FR training samples are already aligned to a pre-defined template, *I.e.*, the positions of landmarks in \mathbf{x} are already available. Furthermore, the final landmark location after applying the transformation can be obtained using Equation 4. Therefore, the proposed constraint does not need landmark estimation and \mathcal{S} can be effectively computed using the FR alignment template, *I.e.*, common across available large-scale datasets, and Equation 4. Section 1 of Supplementary material provides a detailed explanation on \mathcal{S} .

Furthermore, to inject uncertainty to the crafted T_θ , the PGD step size α is randomly sampled from the Gaussian distribution $N(\mu, \sigma^2)$:

$$\theta_{(k+1)} \leftarrow \text{proj}_{\mathcal{S}} (\theta_{(k)} + \alpha \text{sgn} (\nabla_{\theta_{(k)}} L(F_\psi; T_{\theta_{(k)}}(\mathbf{x}), y))); \quad \text{s.t. } \alpha \sim N(\mu, \sigma^2), \quad (9)$$

randomized α significantly helps the ARoFace to increase the diversity of crafted transformations. This diversity is essential to enrich the training with samples expressing the uncertainty inherent to FAE. Following obtaining the adversarial transformation parameter using Equation 9, the F_ψ is trained using a batch containing benign and adversarial samples, minimization of Equation 7. This framework enriches the training by spatial transformations that express FAE.

Table 1: Comparison with SOTA methods on the IJB-B, IJB-C, and TinyFace when the backbone is ResNet-100. ‘*’ indicates re-runs with official code due to missing trained models on the official repository.

Method	Venue	Train Set	IJB-B					IJB-C					TinyFace	
			TAR@FAR		Identification			TAR@FAR		Identification			Rank1	Rank5
			1e-6	1e-5	1e-4	Rank1	Rank5	1e-6	1e-5	1e-4	Rank1	Rank5	Rank1	Rank5
MagFace [46]	CVPR2021	MS1MV2	42.32	90.36	94.51	-	-	89.26	93.67	95.81	-	-	-	-
MagFace+IIC [27]	ICLR2024	MS1MV2	-	-	-	-	-	89.38	93.95	95.89	-	-	-	-
ArcFace [11]	CVPR2019	MS1MV2	38.68	88.50	94.09	-	-	85.66	92.69	95.74	-	-	-	-
ArcFace+CFSM* [41]	ECCV2022	MS1MV2	47.27	90.52	95.21	95.00	96.67	90.83	94.72	96.60	96.19	97.30	64.69	68.80
ArcFace+ARoFace		MS1MV2	48.70	90.83	95.38	95.13	96.71	89.28	94.74	96.66	96.20	97.29	67.32	72.45
ArcFace+VPL [13]	CVPR2021	MS1MV3	-	-	95.56	-	-	-	-	96.76	-	-	-	-
ArcFace+SC [9]	ECCV2020	MS1MV3	-	-	95.25	-	-	-	-	96.61	-	-	-	-
ArcFace [11]	CVPR2019	MS1MV3	40.27	92.09	95.47	95.29	97.01	90.99	95.31	96.81	96.61	97.66	63.81	68.05
ArcFace+ARoFace		MS1MV3	42.31	92.85	95.68	95.79	97.44	91.48	95.69	96.87	97.07	98.09	67.54	71.05
ArcFace* [11]	CVPR2019	WebFace4M	44.67	92.24	95.76	96.17	97.51	91.45	95.43	97.16	97.42	98.24	71.11	74.38
ArcFace+ARoFace		WebFace4M	45.10	92.33	95.83	96.13	97.62	91.48	95.69	96.87	97.51	98.35	73.80	76.53
AdaFace [36]	CVPR2022	WebFace4M	44.48	92.26	96.03	96.26	97.68	90.43	95.34	97.39	97.52	98.35	72.02	74.52
AdaFace+ARoFace		WebFace4M	44.21	93.57	96.35	96.23	97.71	91.57	95.98	97.51	97.59	98.41	73.98	76.47
AdaFace [36]	CVPR2022	WebFace12M	47.49	93.13	96.30	96.28	97.72	89.47	95.94	97.54	97.56	98.38	72.29	74.97
AdaFace+ARoFace		WebFace12M	44.76	93.14	96.23	96.23	97.93	89.80	96.20	97.60	97.61	98.57	74.00	76.87

Consequently, it improves the model robustness to FAE which eventually enhances the FR performance on LQ faces. Our proposal is orthogonal to the FR training objective and improves LQ FR without facing the issue of identity preservation, and reliance on a *a priori* defined target dataset. The whole process is shown in Figure 2 and Algorithm 1.

4 Experiments

4.1 Dataset

We utilize the cleaned version of MS-Celeb-1M [21], provided by [11, 14] as our training dataset, *i.e.*, MS1MV3. This dataset consists of almost 5M images from 90K identities. As per the conventional FR framework, all used datasets in our work are aligned and transformed to 112×112 pixels. We evaluate ARoFace on TinyFace [6], IJB-B [69], IJB-C [45], and IJB-S [32] datasets.

TinyFace [6] is an FR evaluation dataset comprising 169,403 LQ face images across 5,139 identities, designed for 1: N recognition. The average image size of this dataset is 20 by 16 pixels.

IJB-B and IJB-C. IJB-B [69] contains approximately 21.8K images (11.8K faces and 10K non-face) and 7k videos (55K frames), representing a total of 1,845 identities. IJB-C [45], an extension of IJB-B, includes 31.3K images and 117.5K frames from 3,531 identities. IJB-B and IJB-C contain both HQ and LQ samples and have been widely used for evaluating the FR models [36].

IJB-S [32] is recognized as one of the most challenging FR benchmarks, primarily utilizing samples from real-world surveillance videos, see Figure 3a. This dataset comprises 350 surveillance videos, totaling 30 hours, from 202 identities. Additionally, there are seven HQ photos for each subject. This dataset is characterized by three keywords:

- **Surveillance:** Refers to the use of surveillance video footage.

Table 2: Comparison with SOTA methods on the IJB-S benchmark when the backbone is ResNet-100. ARoFace outperforms all the baselines by a considerable margin.

Method	Venue	Dataset	Surveillance-to-Single			Surveillance-to-Booking			Surveillance-to-Surveillance		
			Rank1	Rank5	1	Rank1	Rank5	1	Rank1	Rank5	1
ArcFace [11]	CVPR2019	MS1MV2	57.35	64.42	41.85	57.36	64.95	41.23	-	-	-
PFE [53]	ICCV2019	MS1MV2	50.16	58.33	31.88	53.60	61.75	35.99	9.20	20.82	0.84
URL [55]	ICCV2020	MS1MV2	59.79	65.78	41.06	61.98	67.12	42.73	-	-	-
ArcFace+ARoFace		MS1MV2	61.65	67.6	47.87	60.66	67.33	46.34	18.31	32.07	2.23
ArcFace [11]	CVPR2019	WebFace4M	69.26	74.31	57.06	70.31	75.15	56.89	32.13	46.67	5.32
ArcFace+ARoFace		WebFace4M	70.96	75.54	58.67	71.70	75.24	58.06	32.95	50.30	6.81
AdaFace [36]	CVPR2022	WebFace12M	71.35	76.24	59.40	71.93	76.56	59.37	36.71	50.03	4.62
AdaFace+ARoFace		WebFace12M	72.28	77.93	61.43	73.01	79.11	60.02	40.51	50.90	6.37

- **Single:** Denotes the utilization of a single HQ enrollment image.
- **Booking:** Multiple enrollment images taken from different viewpoints.

4.2 Implementation Details

We adopt a modified version of ResNet-100 [11] as our backbone. The training is done for 28 epochs using ArcFace loss, with exceptions noted where applicable. We employ SGD as the optimizer, with a cosine annealing learning rate starting from 0.1, a weight decay of 0.0001, and a momentum of 0.9. The component of θ governing scale transformation is initialized with $N(\mu = 1, \sigma^2 = 0.01)$, while parameters for rotation and translation are derived from $N(\mu = 0, \sigma^2 = 0.01)$. Furthermore, α is sampled from $N(\mu = 0, \sigma^2 = 0.01)$. Section 2 of Supplementary Material provides detailed experiments on these parameters. During training, each GPU handles a mini-batch of size 512, utilizing four Nvidia RTX 6000 GPUs. For a fair comparison, when a checkpoint for a specific method was not available, we used the official code released by the authors and the optimal hyperparameters as recommended in their publication for reproducing their results.

4.3 Comparison with SOTA Methods

Table 1 compares ARoFace performance on IJB-B, IJB-C, and TinyFace against SOTA methods. These results demonstrate that ARoFace consistently sets new SOTA benchmarks across a variety of metrics and datasets. These advancements using different training datasets suggest that ARoFace’s benefits are not confined to any specific training dataset. When employing the WebFace4M, MS1MV3, and MS1MV2 training sets, ARoFace outperforms its competitors by 0.43%, 2.04%, and 1.43% at FAR= $1e-6$ on IJB-B, respectively. Concretely, on IJB-C, ARoFace surpasses AdaFace by 1.14% using WebFace4M and ArcFace by 0.49% using MS1MV3 at FAR= $1e-6$. Using MS1MV2 as the training data, ARoFace (ArcFace+ARoFace) outperforms CFSSM [41] (ArcFace+CFSSM) by 2.63%, and 3.65% improvements in Rank1, and Rank5 TinyFace identification, respectively. These enhancements over CFSSM, which prioritizes the inclusion of synthesized LQ data in the training, underscore the significance of mitigating the susceptibility to FAE over the introduction of LQ instances to the training.

Table 3: Comparison between baseline and baseline+ARoFace on IJB-B and IJB-C dataset on aligned and alignment-perturbed inputs.

Method	Input Alignment	IJB-B				IJB-C			
		1e-5	1e-4	1e-3	1e-2	1e-5	1e-4	1e-3	1e-2
ArcFace [11]	Original	92.09	95.47	96.92	97.84	95.31	96.81	97.88	98.54
ArcFace+ARoFace	Original	92.85	95.68	96.97	97.88	95.69	96.87	97.95	98.56
General improvement		+0.76	+0.21	+0.05	+0.04	+0.38	+0.06	+0.07	+0.02
ArcFace [11]	Perturbed	45.74	83.64	88.80	92.34	16.17	68.89	85.53	91.11
ArcFace+ARoFace	Perturbed	61.09	90.85	93.17	94.86	28.92	84.50	91.40	96.27
Improvement to misalignment		+15.35	+7.21	+4.37	+2.52	+12.75	+15.61	+5.87	+5.16
AdaFace [36]	Original	90.86	95.84	97.34	98.31	95.33	97.09	98.14	98.89
AdaFace+ARoFace	Original	90.05	95.91	97.61	98.34	95.31	97.12	98.24	98.91
General improvement		-0.81	+0.04	+0.27	+0.03	-0.02	+0.03	+0.1	+0.03
AdaFace [36]	Perturbed	15.71	54.88	84.44	92.46	2.16	19.12	57.69	88.14
AdaFace+ARoFace	Perturbed	48.61	82.44	89.40	92.69	20.19	64.52	85.23	91.61
Improvement to misalignment		+32.9	+27.56	+4.96	+0.23	+18.03	+45.4	+27.54	+3.47

Using WebFace4M as the training data, ARoFace boosts AdaFace performance in TinyFace evaluation by 1.82% and 2.11% in Rank1 and Rank5, respectively. Furthermore, using WebFace12M as the training data, ARoFace enhances AdaFace by 2.35%, and 2.5% in Rank1 and Rank5, respectively. These consistent enhancements demonstrate that ARoFace maintains its effectiveness as dataset size increases, *i.e.*, from 4M to 12M. Remarkably, ARoFace achieves top performance on TinyFace without lowering the results on IJB-B and IJB-C benchmarks, outperforming its competitors in most verification and identification metrics. IJB-B and IJB-C datasets comprise both HQ and LQ faces [36], as depicted in Figure 3b. Hence, consistent improvements across IJB-B, IJB-C, and TinyFace show ARoFace generalization and its ability to enhance LQ FR without compromising performance on HQ faces.

Furthermore, Table 2 shows ARoFace’s performance on the IJB-S dataset against its competitors. ARoFace establishes a new SOTA performance by significantly outperforming previous baselines. In particular, ARoFace exceeds the ArcFace baseline with the MS1MV2 training set, marking improvements of 6.02% in Surveillance-to-Single and 5.11% in Surveillance-to-Booking verification. This notable progress not only underscores the face alignment challenges in LQ benchmarks like IJB-S but also demonstrates our method’s ability to improve FR model robustness against FAE. The consistent improvement across IJB-S metrics highlights the crucial role of FAE in lowering the discriminative power of existing SOTA FR models, such as ArcFace and AdaFace. These enhancements, observed with different training datasets, affirm that our method’s effectiveness is independent of the training set.

4.4 Analysis on Robustness to Face Alignment Error

Face alignment is an essential prerequisite for almost all available FR models. In Figure 1, we empirically illustrate the severe vulnerability of SOTA FR networks to even mild FAE in IJB-B and IJB-C. Here, we further examine the effectiveness of ARoFace in improving model robustness against FAE. Table 3 compares

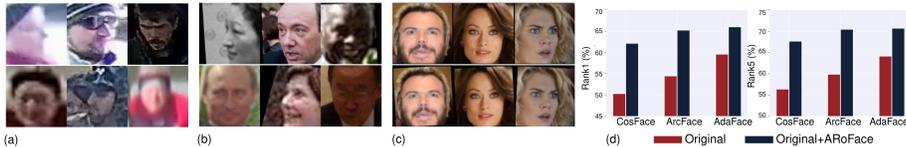


Fig. 3: (a, b) Visualizing samples from IJB-S and IJB-B datasets respectively. IJB-B consists of both HQ and LQ instances while IJB-S only consists of LQ probe instances. (c) Visualizing the benign (top : \mathbf{x}) and their corresponding adversarial example (bottom : $T_\theta(\mathbf{x})$) produced by ARoFace. (d) Orthogonality of the ARoFace to different FR objective functions. In all scenarios, integrating ARoFace into training significantly improved performance on TinyFace.

the performance of ArcFace, and AdaFace under their original training regimes and our modified training approach. Introducing ARoFace results in considerable improvement in the robustness to FAE for both methods. Specifically, the ArcFace performance on the perturbed images increases by 15.35% and 12.75% in $\text{TAR@FAR}=1e-5$, and 7.21% and 15.61% in $\text{TAR@FAR}=1e-4$ for IJB-B and IJB-C, respectively. Consistently, the Adaface performance on the perturbed images improved by 32.9% and 18.03% in $\text{TAR@FAR}=1e-5$, and 27.56% and 45.4% in $\text{TAR@FAR}=1e-4$ for IJB-B and IJB-C, respectively.

These significant improvements showcase the efficacy of the proposed method in making the FR network more robust against FAE. Note that the inclusion of classical augmentation techniques in the training of FR models results in decreasing model generalization [36, 41, 51]. Notably, integrating ARoFace resulted in performance improvement in aligned and alignment perturbed inputs. This consistent improvement is particularly noteworthy which emphasizes the generalizability of the proposed method and highlights that ARoFace does not sacrifice the performance on aligned samples to gain robustness against FAE.

4.5 Orthogonal Improvement to Angular Margin

ARoFace aims to improve LQ FR by injecting misaligned samples into the training. Figure 3c shows original and augmented samples along each other. Here, we evaluate the effectiveness of ARoFace across different angular margins, depicted in Figure 3d. Specifically, we modified the training code of three SOTA methods, *i.e.*, CosFace, ArcFace, and AdaFace, by incorporating our training policy and using the CASIA-WebFace [81] as the training data. These results suggest the orthogonality of ARoFace to existing angular penalty losses. Thus, ARoFace can be seamlessly integrated into SOTA FR frameworks as a plug-and-play module, increasing their robustness to FAE and ultimately improving performance, particularly in LQ evaluations.

4.6 Runtime Overhead

One significant benefit of the proposed method over [41, 54] is its ability to bypass the need for complex image generation procedure, which requires two-step optimization, *i.e.*, training a generator and then integrating it to the FR

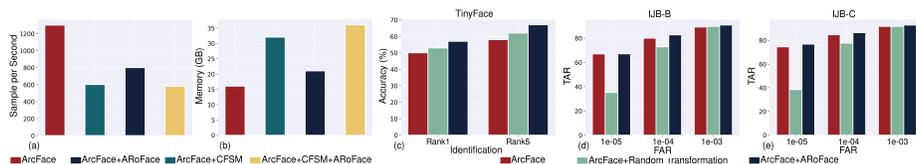


Fig. 4: (a, b) Training speed and GPU memory consumption comparison between CFSM and ARoFace: ARoFace significantly enhances training efficiency and reduces GPU memory consumption compared to CFSM. (c, d, e) Comparing the evaluation performance between employing adversarial vs. random spatial transformation during training: Adversarial improves performance, while random fails on IJB-B and IJB-C.

training. Our proposal crafts an affine transformation [28] to apply directly to the input image, thereby eliminating the need for two-step optimization.

Moreover, ARoFace employs only four additional trainable parameters, responsible for a negligible increase in computational load. The majority of the overhead computation time and memory consumption is due to adversarial training. In Figure 4a and b, we compare our method’s training speed and GPU memory usage with CFSM of Liu *et al.* [41], which also aims to enhance LQ FR. By avoiding the usage of image generation models, *i.e.*, GAN in [41], we have increased CFSM’s [41] training speed by 25% and decreased the memory consumption by 35%. These significant improvements in training speed and reduction in memory consumption, coupled with improved performance as shown in Tables 1 and 2, emphasize the capability of our method for large-scale FR.

Furthermore, we also compare (ArcFace+CFSM+ARoFace) with the original CFSM setup (ArcFace+CFSM). CFSM already employs adversarial optimization. Hence, the integration of ARoFace into CFSM results in a negligible decrease in training speed and an increase in memory usage, as shown in Figure 4a, b. The extra computational effort from integrating ARoFace is minimal, confirming that extra computational load stems from the adversarial training which is necessary for integrating augmentation into FR training [41].

4.7 Ablation on Transformation Components

In Section 3.3, we note that the spatial transformation employed in ARoFace comprises scaling, rotation, and translation. Here, we investigate the effect of each component on the performance of ARoFace, *i.e.*, how much each transformation component contributes to the performance enhancement. To this end, we conduct an ablation study on TinyFace with ArcFace serving as the loss function and CASIA-WebFace [81] as the training dataset. These experiments, as shown in Table 4, show that scaling is the most beneficial transformation, enhancing performance by nearly 5%. This aligns with previous re-

Table 4: Ablation on transformation components when the training data is CASIA-WebFace, backbone is ResNet-100 and the objective function is ArcFace.

Scale	Rotation	Translation	Rank1	Rank5
×	×	×	54.47	59.89
✓	×	×	59.70	65.17
×	✓	×	57.32	61.93
×	×	✓	57.98	62.57
✓	✓	×	62.32	67.01
✓	✓	✓	65.26	70.46

search [36, 41, 54], highlighting scaling’s role in generating LQ inputs. Translation and rotation each contribute approximately 3% improvement, reflecting their similar influence on face alignment.

4.8 Random vs. Adversarial

Liu *et al.* [41] showed that the performance boost is marginal without adversarial optimization, *I.e.*, randomly generate LQ samples. Here, we investigate this effect in the proposed framework in Figure 4c, d, and e. The results indicate that employing spatial augmentation, either random or with an adversarial signal, increases the baseline identification performance on LR inputs, *I.e.*, TinyFace. However, random augmentation destroys the discriminative power of the FR model in HQ samples, *I.e.*, a drastic decrease in IJB-B and IJB-C datasets. The key benefit of employing differentiable spatial transformation of ARoFace is the performance gain across datasets with different image characteristics, *I.e.*, IJB-B, IJB-C, IJB-S, and TinyFace, and not sacrificing the performance on aligned samples to gain robustness to FAE.

4.9 Effect of Random Sampling of α

Table 5 compares the ARoFace performance with and without employing randomly sampled α . In these experiments, when α is fixed, its value equals the mean of the distribution of random, *I.e.*, $\alpha_{random} \sim N(\mu, \sigma^2)$, and $\alpha_{fixed} = \mu$. In TinyFace, employing α_{random} provides notable improvement over α_{fixed} . This significant improvement highlights the fact that the FAE in LQ images is notably diverse and adding uncertainty through α_{random} helps ARoFace to increase the robustness to FAE. Furthermore, the performance on IJB-B and IJB-C is almost the same in both scenarios. We attribute this to the fact that these datasets are a mixture of LQ and HQ faces.

Table 5: Experiment on α . Verification performance TAR@FAR=1e-4 is reported for IJB-B and IJB-C. Training data is CASIA-WebFace, backbone is ResNet-100 and ArcFace is the objective function.

	TinyFace		IJB-B	IJB-C
	Rank1	Rank5		
α_{fixed}	62.15	67.48	58.36	52.95
α_{random}	65.26	70.46	58.47	52.88

5 Conclusion

In this paper, based on our observation on the susceptibility of current FR models to FAE in LQ faces, we proposed to consider FAE as another image degradation factor that is specifically tailored to FR. Our method employs a differentiable spatial transformation combined with adversarial data augmentation to craft samples expressing FAE and add them to the training of the FR model. This framework, allows the FR model to become exposed to FAE during the training and gain robustness against FAE. The efficacy of the proposed method is evaluated through various experiments and evaluation across different benchmarks, including IJB-B, IJB-C, IJB-S, and TinyFace.

References

1. Antoniou, A., Storkey, A., Edwards, H.: Data augmentation generative adversarial networks. arXiv preprint arXiv:1711.04340 (2017)
2. Boutros, F., Fang, M., Klemt, M., Fu, B., Damer, N.: Cr-fiq: Face image quality assessment by learning sample relative classifiability. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5836–5845 (June 2023)
3. Chai, J.C.L., Ng, T.S., Low, C.Y., Park, J., Teoh, A.B.J.: Recognizability embedding enhancement for very low-resolution face recognition and quality estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9957–9967 (2023)
4. Chen, Y., Tai, Y., Liu, X., Shen, C., Yang, J.: Fsrnet: End-to-end learning face super-resolution with facial priors. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2492–2501 (2018)
5. Cheng, Z., Zhu, X., Gong, S.: Low-resolution face recognition. In: Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14. pp. 605–621. Springer (2019)
6. Cheng, Z., Zhu, X., Gong, S.: Low-resolution face recognition. In: Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14. pp. 605–621. Springer (2019)
7. Dabouei, A., Soleymani, S., Dawson, J., Nasrabadi, N.: Fast geometrically-perturbed adversarial faces. In: 2019 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1979–1988. IEEE (2019)
8. Deb, D., Zhang, J., Jain, A.K.: Advfaces: Adversarial face synthesis. In: 2020 IEEE International Joint Conference on Biometrics (IJCB). pp. 1–10. IEEE (2020)
9. Deng, J., Guo, J., Liu, T., Gong, M., Zafeiriou, S.: Sub-center arcface: Boosting face recognition by large-scale noisy web faces. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16. pp. 741–757. Springer (2020)
10. Deng, J., Guo, J., Ververas, E., Kotsia, I., Zafeiriou, S.: Retinaface: Single-shot multi-level face localisation in the wild. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5203–5212 (2020)
11. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4690–4699 (2019)
12. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2019)
13. Deng, J., Guo, J., Yang, J., Lattas, A., Zafeiriou, S.: Variational prototype learning for deep face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11906–11915 (2021)
14. Deng, J., Guo, J., Zhang, D., Deng, Y., Lu, X., Shi, S.: Lightweight face recognition challenge. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019)
15. Dong, Y., Liao, F., Pang, T., Su, H., Zhu, J., Hu, X., Li, J.: Boosting adversarial attacks with momentum. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 9185–9193 (2018)

16. Dong, Y., Su, H., Wu, B., Li, Z., Liu, W., Zhang, T., Zhu, J.: Efficient decision-based black-box adversarial attacks on face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7714–7722 (2019)
17. Engstrom, L., Tran, B., Tsipras, D., Schmidt, L., Madry, A.: Exploring the landscape of spatial robustness. In: International conference on machine learning. pp. 1802–1811. PMLR (2019)
18. Ge, S., Zhao, S., Li, C., Zhang, Y., Li, J.: Efficient low-resolution face recognition via bridge distillation. *IEEE Transactions on Image Processing* **29**, 6898–6908 (2020)
19. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint arXiv:1412.6572 (2014)
20. Gu, Y., Wang, X., Xie, L., Dong, C., Li, G., Shan, Y., Cheng, M.M.: Vqfr: Blind face restoration with vector-quantized dictionary and parallel decoder. In: ECCV (2022)
21. Guo, Y., Zhang, L., Hu, Y., He, X., Gao, J.: Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In: Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14. pp. 87–102. Springer (2016)
22. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
23. Hennings-Yeomans, P.H., Baker, S., Kumar, B.V.: Simultaneous super-resolution and feature extraction for recognition of low-resolution faces. In: 2008 IEEE Conference on computer vision and pattern recognition. pp. 1–8. IEEE (2008)
24. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: CVPR (2017)
25. Huang, G.B., Mattar, M., Berg, T., Learned-Miller, E.: Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In: Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition (2008)
26. Huang, R., Zhang, S., Li, T., He, R.: Beyond face rotation: Global and local perception gan for photorealistic and identity preserving frontal view synthesis. In: Proceedings of the IEEE international conference on computer vision. pp. 2439–2448 (2017)
27. Huang, Y., Wang, Y., Yang, L., Wang, L.: Enhanced face recognition using intra-class incoherence constraint. In: The Twelfth International Conference on Learning Representations (2024), <https://openreview.net/forum?id=uELjxVbrqG>
28. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. *Advances in neural information processing systems* **28** (2015)
29. Jo, Y., Yang, S., Kim, S.J.: Investigating loss functions for extreme super-resolution. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops. pp. 424–425 (2020)
30. Kalarot, R., Li, T., Porikli, F.: Component attention guided face super-resolution network: Cagface. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 370–380 (2020)
31. Kalka, N.D., Maze, B., Duncan, J.A., O'Connor, K., Elliott, S., Hebert, K., Bryan, J., Jain, A.K.: Ijb-s: Iarpa janus surveillance video benchmark. In: 2018 IEEE 9th international conference on biometrics theory, applications and systems (BTAS). pp. 1–9. IEEE (2018)

32. Kalka, N.D., Maze, B., Duncan, J.A., O'Connor, K., Elliott, S., Hebert, K., Bryan, J., Jain, A.K.: Ijb-s: Iarpa janus surveillance video benchmark. In: 2018 IEEE 9th international conference on biometrics theory, applications and systems (BTAS). pp. 1–9. IEEE (2018)
33. Kanazawa, A., Jacobs, D.W., Chandraker, M.: Warpnet: Weakly supervised matching for single-view reconstruction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3253–3261 (2016)
34. Kim, D., Kim, M., Kwon, G., Kim, D.S.: Progressive face super-resolution via attention to facial landmark. arXiv preprint arXiv:1908.08239 (2019)
35. Kim, H.I., Yun, K., Ro, Y.M.: Face shape-guided deep feature alignment for face recognition robust to face misalignment. *IEEE Transactions on Biometrics, Behavior, and Identity Science* **4**(4), 556–569 (2022)
36. Kim, M., Jain, A.K., Liu, X.: Adaface: Quality adaptive margin for face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18750–18759 (2022)
37. Komkov, S., Petiushko, A.: Advhat: Real-world adversarial attack on arcface face id system. In: 2020 25th International Conference on Pattern Recognition (ICPR). pp. 819–826. IEEE (2021)
38. Kurakin, A., Goodfellow, I., Bengio, S.: Adversarial machine learning at scale. arXiv preprint arXiv:1611.01236 (2016)
39. Li, J., Guo, Z., Li, H., Han, S., Baek, J.w., Yang, M., Yang, R., Suh, S.: Rethinking feature-based knowledge distillation for face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20156–20165 (2023)
40. Li, X., Li, W., Ren, D., Zhang, H., Wang, M., Zuo, W.: Enhanced blind face restoration with multi-exemplar images and adaptive spatial feature fusion. *CVPR* (2020)
41. Liu, F., Kim, M., Jain, A., Liu, X.: Controllable and guided face synthesis for unconstrained face recognition. In: Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XII. pp. 701–719. Springer (2022)
42. Liu, W., Wen, Y., Yu, Z., Li, M., Raj, B., Song, L.: Sphereface: Deep hypersphere embedding for face recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 212–220 (2017)
43. Ma, C., Jiang, Z., Rao, Y., Lu, J., Zhou, J.: Deep face super-resolution with iterative collaboration between attentive recovery and landmark estimation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5569–5578 (2020)
44. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. arXiv preprint arXiv:1706.06083 (2017)
45. Maze, B., Adams, J., Duncan, J.A., Kalka, N., Miller, T., Otto, C., Jain, A.K., Niggel, W.T., Anderson, J., Cheney, J., et al.: Iarpa janus benchmark-c: Face dataset and protocol. In: 2018 international conference on biometrics (ICB). pp. 158–165. IEEE (2018)
46. Meng, Q., Zhao, S., Huang, Z., Zhou, F.: Magface: A universal representation for face recognition and quality assessment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14225–14234 (2021)
47. Nair, N.G., Mei, K., Patel, V.M.: At-ddpm: Restoring faces degraded by atmospheric turbulence using denoising diffusion probabilistic models. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3434–3443 (2023)

48. Niemeijer, J., Schwonberg, M., Termöhlen, J.A., Schmidt, N.M., Fingscheidt, T.: Generalization by adaptation: Diffusion-based domain extension for domain-generalized semantic segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2830–2840 (2024)
49. Peng, X., Tang, Z., Yang, F., Feris, R.S., Metaxas, D.: Jointly optimize data augmentation and network training: Adversarial data augmentation in human pose estimation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2226–2234 (2018)
50. Robbins, W., Boulton, T.E.: On the effect of atmospheric turbulence in the feature space of deep face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1618–1626 (2022)
51. Saadabadi, M.S.E., Malakshan, S.R., Zafari, A., Mostofa, M., Nasrabadi, N.M.: A quality aware sample-to-sample comparison for face recognition. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 6129–6138 (2023)
52. Sharif, M., Bhagavatula, S., Bauer, L., Reiter, M.K.: Accessorize to a crime: Real and stealthy attacks on state-of-the-art face recognition. In: Proceedings of the 2016 acm sigsac conference on computer and communications security. pp. 1528–1540 (2016)
53. Shi, Y., Jain, A.K.: Probabilistic face embeddings. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6902–6911 (2019)
54. Shi, Y., Jain, A.K.: Boosting unconstrained face recognition with auxiliary unlabeled data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2795–2804 (2021)
55. Shi, Y., Yu, X., Sohn, K., Chandraker, M., Jain, A.K.: Towards universal representation learning for deep face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6817–6826 (2020)
56. Shin, S., Lee, J., Lee, J., Yu, Y., Lee, K.: Teaching where to look: Attention similarity knowledge distillation for low resolution face recognition. In: Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XII. pp. 631–647. Springer Nature Switzerland Cham (2022)
57. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
58. Singh, M., Nagpal, S., Singh, R., Vatsa, M.: Dual directed capsule network for very low resolution image recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 340–349 (2019)
59. Soundararajan, R., Biswas, S.: Machine vision quality assessment for robust face detection. *Signal Processing: Image Communication* **72**, 92–104 (2019)
60. Terhörst, P., Ihlefeld, M., Huber, M., Damer, N., Kirchbuchner, F., Raja, K., Kuijper, A.: Qmagface: Simple and accurate quality-aware face recognition. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3484–3494 (2023)
61. Wang, F., Xiang, X., Cheng, J., Yuille, A.L.: Normface: L2 hypersphere embedding for face verification. In: Proceedings of the 25th ACM international conference on Multimedia. pp. 1041–1049 (2017)
62. Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., Liu, W.: CosFace: Large margin cosine loss for deep face recognition. In: CVPR (2018)
63. Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., Liu, W.: CosFace: Large Margin Cosine Loss for Deep Face Recognition. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition pp. 5265–5274 (2018)

64. Wang, K., Wang, S., Zhang, P., Zhou, Z., Zhu, Z., Wang, X., Peng, X., Sun, B., Li, H., You, Y.: An efficient training approach for very large scale face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4083–4092 (2022)
65. Wang, M., Liu, R., Hajime, N., Narishige, A., Uchida, H., Matsunami, T.: Improved knowledge distillation for training fast low resolution face recognition model. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019)
66. Wang, Y., Pan, X., Song, S., Zhang, H., Huang, G., Wu, C.: Implicit semantic data augmentation for deep networks. *Advances in Neural Information Processing Systems* **32** (2019)
67. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
68. Wang, Z., Zhang, J., Chen, R., Wang, W., Luo, P.: Restoreformer: High-quality blind face restoration from undegraded key-value pairs. *CVPR* (2022)
69. Whitelam, C., Taborsky, E., Blanton, A., Maze, B., Adams, J., Miller, T., Kalka, N., Jain, A.K., Duncan, J.A., Allen, K., et al.: Iarpa janus benchmark-b face dataset. In: proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 90–98 (2017)
70. Wu, W., Kan, M., Liu, X., Yang, Y., Shan, S., Chen, X.: Recursive spatial transformer (rest) for alignment-free face recognition. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3772–3780 (2017)
71. Xiao, C., Zhu, J.Y., Li, B., He, W., Liu, M., Song, D.: Spatially transformed adversarial examples. *arXiv preprint arXiv:1801.02612* (2018)
72. Xie, C., Tan, M., Gong, B., Wang, J., Yuille, A.L., Le, Q.V.: Adversarial examples improve image recognition. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 819–828 (2020)
73. Xie, C., Wu, Y., Maaten, L.v.d., Yuille, A.L., He, K.: Feature denoising for improving adversarial robustness. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 501–509 (2019)
74. Xu, X., Meng, Q., Qin, Y., Guo, J., Zhao, C., Zhou, F., Lei, Z.: Searching for alignment in face recognition. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 3065–3073 (2021)
75. Yang, C.Y., Ma, C., Yang, M.H.: Single-image super-resolution: A benchmark. In: Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13. pp. 372–386. Springer (2014)
76. Yang, K., Sun, Y., Su, J., He, F., Tian, X., Huang, F., Zhou, T., Tao, D.: Adversarial auto-augment with label preservation: A representation learning principle guided approach. *Advances in Neural Information Processing Systems* **35**, 22035–22048 (2022)
77. Yang, R., Laurel, J., Misailovic, S., Singh, G.: Provable defense against geometric transformations. In: The Eleventh International Conference on Learning Representations (2022)
78. Yang, T., Ren, P., Xie, X., Zhang, L.: Gan prior embedded network for blind face restoration in the wild. *CVPR* (2021)
79. Yang, T.Y., Chen, Y.T., Lin, Y.Y., Chuang, Y.Y.: Fsa-net: Learning fine-grained structure aggregation for head pose estimation from a single image. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1087–1096 (2019)

80. Yang, W., Zhang, X., Tian, Y., Wang, W., Xue, J.H., Liao, Q.: Deep learning for single image super-resolution: A brief review. *IEEE Transactions on Multimedia* **21**(12), 3106–3121 (2019)
81. Yi, D., Lei, Z., Liao, S., Li, S.Z.: Learning face representation from scratch. *arXiv preprint arXiv:1411.7923* (2014)
82. Yin, B., Wang, W., Yao, T., Guo, J., Kong, Z., Ding, S., Li, J., Liu, C.: Adv-makeup: A new imperceptible and transferable attack on face recognition. *arXiv preprint arXiv:2105.03162* (2021)
83. Yu, X., Fernando, B., Ghanem, B., Porikli, F., Hartley, R.: Face super-resolution guided by facial component heatmaps. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 217–233 (2018)
84. Zangeneh, E., Rahmati, M., Mohsenzadeh, Y.: Low resolution face recognition using a two-branch deep convolutional neural network architecture. *Expert Systems with Applications* **139**, 112854 (2020)
85. Zhang, K., Zhang, Z., Li, Z., Qiao, Y.: Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE signal processing letters* **23**(10), 1499–1503 (2016)
86. Zhao, H., Ying, X., Shi, Y., Tong, X., Wen, J., Zha, H.: Rdcface: radial distortion correction for face recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7721–7730 (2020)
87. Zhao, Y., Zhong, Z., Zhao, N., Sebe, N., Lee, G.H.: Style-hallucinated dual consistency learning for domain generalized semantic segmentation. In: *European Conference on Computer Vision*. pp. 535–552. Springer Nature Switzerland Cham (2022)
88. Zhong, Y., Chen, J., Huang, B.: Toward end-to-end face recognition through alignment learning. *IEEE signal processing letters* **24**(8), 1213–1217 (2017)
89. Zhou, E., Cao, Z., Sun, J.: Gridface: Face rectification via learning local homography transformations. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 3–19 (2018)
90. Zhou, E., Cao, Z., Sun, J.: Gridface: Face rectification via learning local homography transformations. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 3–19 (2018)
91. Zhou, T., Tulsiani, S., Sun, W., Malik, J., Efros, A.A.: View synthesis by appearance flow. In: *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*. pp. 286–301. Springer (2016)
92. Zhu, Z., Huang, G., Deng, J., Ye, Y., Huang, J., Chen, X., Zhu, J., Yang, T., Lu, J., Du, D., et al.: WebFace260M: A benchmark unveiling the power of million-scale deep face recognition. In: *CVPR* (2021)