

DifFIQA: Face Image Quality Assessment Using Denoising Diffusion Probabilistic Models

Žiga Babnik, Peter Peer, Vitomir Štruc

University of Ljubljana, Kongresni trg 12, SI-1000 Ljubljana, Slovenia

{ziga.babnik, vitomir.struc}@fe.uni-lj.si, peter.peer@fri.uni-lj.si




Figure 1. **High-level idea behind the proposed DifFIQA face image quality assessment (FIQA) approach.** The quality of face images corresponds to a considerable degree to the stability of the respective representations in the embedding space of a given face recognition (FR) model. DifFIQA utilizes a diffusion framework to explore the embedding stability through image perturbations caused by the noising and denoising processes. The intuition behind this approach is that the forward (noising) \mathcal{F}_d and backward (denoising) \mathcal{B}_d diffusion processes lead to larger embedding perturbations for lower-quality images (x^l) compared to facial images of higher quality (x^h). By analyzing the impact of both the forward and backward processes on the representation of a given image, DifFIQA is able to infer the corresponding quality and/or generate (FR model specific) quality rankings, as shown on the right. The figure is best viewed electronically.

Abstract

Modern face recognition (FR) models excel in constrained scenarios, but often suffer from decreased performance when deployed in unconstrained (real-world) environments due to uncertainties surrounding the quality of the captured facial data. Face image quality assessment (FIQA) techniques aim to mitigate these performance degradations by providing FR models with sample-quality predictions that can be used to reject low-quality samples and reduce false match errors. However, despite steady improvements, ensuring reliable quality estimates across facial images with diverse characteristics remains challenging. In this paper, we present a powerful new FIQA approach, named DifFIQA, which relies on denoising diffusion probabilistic models (DDPM) and ensures highly competitive results. The main idea behind the approach is to utilize the forward and backward processes of DDPMs to perturb facial images and quantify the impact of these perturbations on the corresponding image embeddings for quality prediction. Because the diffusion-based perturbations are computationally expensive, we also distill the knowledge encoded in DifFIQA into a regression-based quality predictor, called DifFIQA(R), that balances performance and execution time. We evaluate both models in comprehensive experiments on 7 diverse datasets, with 4 target FR models and against 10 state-of-the-art FIQA techniques with highly encouraging results. The source code is available from: <https://github.com/LSIbabnikz/DifFIQA>.

1. Introduction

State-of-the-art face recognition (FR) models achieve near-perfect results on various benchmarks with high-quality facial images, but still struggle in real-world situations, where the quality of the input samples is frequently unknown [1, 13, 41]. For instance, surveillance, a common application of FR, often involves lower quality samples due to unconstrained and covert capture conditions. In such cases, assessing the quality of the face-image samples is crucial. Low-quality samples can mislead the FR models and cause catastrophic false-match errors, leading to privacy breaches or even monetary loss. By determining the quality of input samples and rejecting or requesting recapture of those below a given quality threshold, the stability, and performance of FR models can typically be improved.

Face Image Quality Assessment (FIQA) methods provide FR methods with a quality estimate for each given face sample. In this context, the term *quality* can refer to either the character, fidelity, or utility of the sample, as defined by ISO/IEC 29794-1 [21]. Similarly to most FIQA research, we focus on the biometric utility of the facial samples, rather than the visual quality (character and fidelity) as perceived by humans [35]. Such image characteristics are commonly evaluated by general-purpose Image Quality Assessment (IQA) techniques. Biometric utility encompasses several unknown aspects of the given face sample, including its visual quality, face-specific information, and

the relative biases inherent to the targeted FR model. It can be interpreted as the usefulness (or fitness) of the sample for the recognition task. Several types of FIQA techniques have been proposed over the years. The largest group focuses on training regression models from calculated pseudo reference quality labels [7, 17, 32, 43], with differences between methods in how they calculate the labels. Other approaches include (unsupervised) analytical methods [3, 28, 39] that use reference-free approaches for quality prediction, and model-based solutions [5, 30, 38] that combine the face recognition and quality assessment tasks. While modern FIQA techniques have demonstrated impressive performance, providing reliable quality predictions for diverse facial characteristics is still a challenging task.

In this paper, we introduce a novel FIQA technique, called DiffFIQA (**Diff**usion-based **F**ace **I**mage **Q**uality **A**ssessment), that leverages the image-generation versatility of modern Denoising Diffusion Probabilistic Models (DDPMs) for face quality assessment and generalizes well across a wide variety of datasets and FR models. As shown in Figure 1, DiffFIQA is based on the following two insights:

- **Perturbation robustness:** Images of higher-quality have stable representations in the embedding space of the given FR model and are less effected by noise perturbations introduced by the forward diffusion process.
- **Reconstruction quality:** High-quality samples are easier to reconstruct from partially corrupted (noisy) data with incomplete identity information and exhibit less disparity between the embeddings of the input and denoised samples than low-quality images.

Based on these observations, DiffFIQA analyzes the embedding stability of the input image by perturbing it through the forward as well as backward diffusion process and infers a quality score from the result. To avoid the computationally expensive backward process and speed up computation, we also distill the DiffFIQA approach into a regression-based model, termed DiffFIQA(R). We evaluate both techniques through extensive experiments over multiple datasets and FR models, and show that both techniques lead to highly competitive results when compared to the state-of-the-art.

2. Related Work

In this section, we briefly review existing FIQA solutions, which can conveniently be partitioned into three main groups: (i) *analytical* techniques, (ii) *regression-based* approaches, and (iii) *model-based* methods.

Analytical methods. The vast majority of methods from this group can be viewed as *specialized* general-purpose IQA techniques that focus on quality predictions defined by (i) selected visual characteristics of faces, such as pose, symmetry or interocular distance, and/or (ii) general visual image properties, such as sharpness, illumination, or noise. An early method from this group was presented by

Raghavendra *et al.* in [33], where a three stage approach combining pose, and image texture components was proposed. Another method by Lijun *et al.* [28] combined several face-image characteristics, including alignment, occlusion and pose, into a pipeline for quality score calculation. Several conceptually similar approaches that exploit different (explicit) visual cues have been presented in the literature over the years [12, 14, 23, 31]. However, the performance of such methods is typically limited, as they focus only on the characteristics of the input samples, with no regard to the utilized FR model. Nevertheless, a new group of analytical methods has recently emerged that incorporates information from both, the input face sample as well as the targeted FR system into the quality estimation process. An example of such an approach was presented by Terh r st *et al.* [39] in the form of the SER-FIQ technique. SER-FIQ calculates a quality score from the embedding variations of a given input face sample, caused by using different configurations of dropout layers. Another method, called FaceQAN by Babnik *et al.* [3], relies on adversarial attacks (which are harder to generate for high quality images) to calculate quality scores. Both of these methods achieve excellent results, but are also comparably computationally demanding, due to their reliance on running several instances of the same sample through the given FR model.

Regression-based methods. FIQA techniques from this group typically train a (quality) regression model using some sort of (pseudo) quality labels. Regression-based methods have received considerable attention over recent years, with most of the research exploring effective mechanisms for generating informative pseudo quality annotations. An early technique from this group, by Best-Rowden and Jain [4], for example, used human raters to annotate the (perceived) quality of facial images, and then trained a quality predictor on the resulting quality labels. Another technique, named FaceQnet [16, 17], relied on embedding comparisons with the highest quality image of each individual to estimate reference quality scores. Here, the highest quality images of each individual were determined using an external quality compliance tool and a ResNet-based regressor was then trained on the extracted quality labels. A more recent approach, called PCNet [43], used a large number of mated image pairs (i.e., a pair of distinct images of the same individual), to train a CNN-based regression model, where the quality labels were defined by the embedding similarity of the mated pairs. The SDD-FIQA approach, by Ou *et al.* [32], extended this concept to also include non-mated (impostor) pairs, (i.e., two unique images of different individuals), where the label for a single image was computed as the Wasserstein distance between the mated and non-mated score distributions. LightQNet, by Chen *et al.* [7], trained a lightweight model, by employing an identification quality loss using quality scores computed from various image comparisons. While regression-based meth-

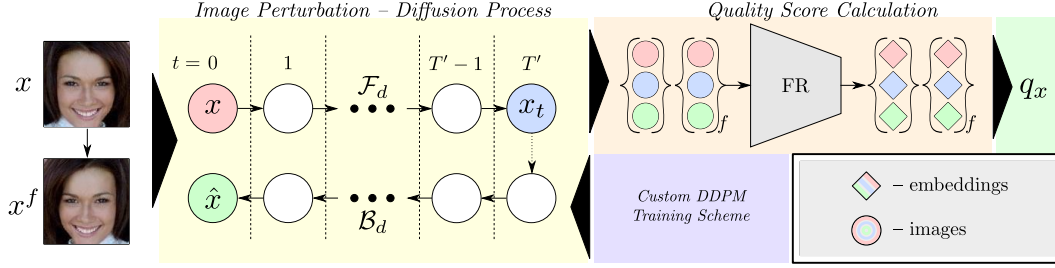


Figure 2. **Overview of DiffFIQA.** The proposed quality assessment pipeline consists of two main parts: the *Diffusion Process* and the *Quality-Score Calculation*. The diffusion process uses an encoder-decoder UNet model (D), trained using an extended DDPM training scheme, that helps to generate higher-quality (restored) images. The custom DDPM model is used in the Diffusion Process, which generates noisy x_t and reconstructed \hat{x} images using the forward and backward diffusion processes, respectively. To capture the effect of facial pose on the quality estimation procedure, the process is repeated with a horizontally flipped image x^f . The Quality Score Calculation part then produces and compares the embeddings of the original images and the images generated by the diffusion part.

ods in general perform well over a variety of benchmarks and state-of-the-art FR models, their main weakness is the lack of specialization. As the optimal quality estimate for a given input image, is by *definition* FR model specific [2, 21], regression-based techniques may require retraining towards the targeted FR model to ensure ideal performance.

Model-based methods. The last group of techniques combines face-image quality assessment and face recognition into a single task. One such technique, PFE by Shi and Jain [38], learned to predict a pair of vectors from the input image, i.e., a mean and a variance vector. The mean vector can be thought of as the embedding of the input sample image, while the variance vector represents the sample’s variability, and can be used to calculate the sample quality. The presented method inspired several new (uncertainty-aware) methods [8, 27, 45], further improving on the performance of PFE. Another notable technique, called MagFace [30], extended the popular ArcFace loss [10] by incorporating a magnitude-aware angular margin term, which dynamically adjusts class boundaries. The embeddings produced by MagFace encode quality in the magnitude of the embedding itself, and can hence be easily inferred. A powerful FIQA technique, called the CR-FIQA, was recently proposed by Boutros *et al.* in [5]. CR-FIQA calculates the quality of the input samples as the ratio between the positive class center and nearest negative class center in a classification task setting, and was demonstrated to produce highly competitive results across various datasets and settings.

Our contribution. The DiffFIQA technique, the *main contribution* of this work, can be seen as an analytical method that relies on the capabilities of a DDPM in combination with a chosen FR model. From a conceptual point of view, it is most closely related to FaceQgen [15], a FIQA technique that uses a (GAN-based) generator model for synthesizing high-quality versions of the input samples and the respective discriminator (that aims to distinguish between genuine high-quality images and poorly restored ones) for quality scoring. Unlike FaceQgen, which analyzes the differences

between the original and restored images independently of the target FR model, DiffFIQA utilizes results from the forward (i.e., noising/degradation) as well as backward (denoising/restoration) diffusion processes and quantifies the embedding variability/uncertainty in the embedding space of a selected FR model for quality estimation. As we show later in the experimental section, this leads to highly competitive FIQA results when compared to the state-of-the-art.

3. Methodology

The stability of the image representations in the embedding space of a given FR model is highly indicative of the input-image quality, as demonstrated by the success of various recent FIQA techniques [3, 39]. One way to explore this stability is by causing perturbations in the image space and analyzing the impact of the perturbations in the embedding space of the targeted FR model. This can, for example, be achieved by using the forward and backward processes of modern diffusion approaches where: the forward process adds some amount of noise to the sample, and the backward process tries to remove the noise, by reconstructing the original. Our main contribution, the DiffFIQA technique, takes advantage of the proposed idea, as illustrated in Figure 2, and employs a custom DDPM model for the generation of noisy and reconstructed images. The generated images are then passed through a chosen FR model to explore the impact of the perturbations on the variability of the embedding corresponding to the input image.

3.1. Preliminaries

To make the paper self-contained, we briefly present the main concept behind denoising diffusion probabilistic models (DDPMs), with a focus on their application within our approach. More information on the theoretical background and applications of diffusion models can be found in [9].

In general, DDPMs represent a special type of generative model that learns to model (image) data distributions through two types of processes: a forward (noising) process

and backward (denoising) process [9, 22]. The **forward diffusion process** \mathcal{F}_d iteratively adds noise to the given input image x_0 , by sampling from a Gaussian distribution $\mathcal{N}(0, I)$. The result of this process is a noisy sample x_t , where t is the number of time steps chosen from the sequence $\{0, 1, \dots, T\}$. The whole forward process \mathcal{F}_d can be presented as a Markov chain given by

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t|x_{t-1}\sqrt{1-\beta_t}, \beta_t I), \quad (1)$$

where β_t is a variance parameter that defines how much noise is added to the sample at the time instance t of the forward process. By making use of the reparameterization trick [18, 25], any sample x_t can be obtained directly from the input sample x_0 , i.e.:

$$q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1-\bar{\alpha}_t)I), \quad (2)$$

where $\bar{\alpha}_t = \prod_{i=0}^t (1 - \beta_i)$.

The **backward diffusion process** \mathcal{B}_d attempts to iteratively denoise the generated samples x_t , using a deep neural network model D_θ parameterized by θ , according to

$$p(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \tilde{\beta}_t I), \quad (3)$$

where $t = T, \dots, 0$, $\tilde{\beta}_t = \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t}\beta_t$, and

$$\mu_\theta(x_t, t) = \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1-\bar{\alpha}_t}x_0 + \frac{\sqrt{\bar{\alpha}_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}x_t. \quad (4)$$

The network is trained to optimize μ_θ , by minimizing the \mathcal{L}_2 loss

$$\mathcal{L}_2 = \mathbb{E}_{t, x_0} \|D_\theta(x_t, t) - x_0\|^2, \quad (5)$$

where $D_\theta(x_t, t)$ is the reconstructed and x_0 the input image. In the remainder of the paper, we drop the subscript θ and use D to denote the deep neural network, which is represented by an unconditional UNet model.

3.2. Overview of DiffIQA

Given a face image x , the goal of DiffIQA is to estimate the quality score $q_x \in \mathbb{R}$, by exploring the effects of the forward and backward diffusion processes of a custom DDPM model D on the image representation e_x in the embedding space of a given FR model M . DiffIQA consists of two main steps, dedicated to: (i) *image perturbation* and (ii) *quality-score calculation*. The image perturbation step uses the forward diffusion process \mathcal{F}_d to create a noisy sample x_t from the input image x and the backward process \mathcal{B}_d to generate the restored (denoised) image \hat{x} . In the quality-score calculation step, the representations $e_x, e_{x_t}, e_{\hat{x}}$ corresponding to the input x , noisy x_t and restored image \hat{x} , are calculated using the FR model M and then analyzed for disparities to infer the final quality score q_x of the input sample x . To also capture pose-related quality information, DiffIQA repeats the entire process using a horizontally flipped version x^f of the input image x , as also illustrated in Figure 2.




Figure 3. **Presentation of the extended DDPM learning scheme.** Given a training sample x and a time step t , the proposed approach generates a time step dependent degraded image x'_t , by combining the original with a degraded image using the function Y . The image x'_t is then used to generate a noisy sample x_t using the standard forward diffusion approach \mathcal{F}_d . A UNet (D) is then trained to reconstruct the input sample in the backward process \mathcal{B}_d .

3.3. Extended DDPM Training

To train the DDPM model D (i.e., a UNet [34]) needed by DiffIQA, we extend the standard training process of diffusion models to incorporate time dependent image degradations, as illustrated in Figure 3. These additional (time-dependent) degradations allow the model to learn to gradually reverse these degradations and, in turn, to construct higher quality images during the backward process. Formally, given an input face image $x_0 = x$, the training procedure first constructs a degraded image $x' = d(x)$, where $d(\cdot)$ is some degradation function. A time step $t \in [0, T]$ is then selected for the given sample, from which a time-dependent degraded image is computed as follows:

$$x'_t = Y(x_0, x', t) = (1 - \ddot{\alpha}_t)x_0 + \ddot{\alpha}_t x' \quad (6)$$

where $\ddot{\alpha}_t$ is calculated as $\sin(\frac{t}{T} \cdot \frac{\pi}{2})$. The degraded image x'_t is then used to produce the noisy sample x_t using (2). Here, $\ddot{\alpha}_t$ is a time-dependent variable that monotonically increases on the interval $t \in [0, T]$, such that $\ddot{\alpha}_{t=0} = 0$ and $\ddot{\alpha}_{t=T} = 1$. In other words, at time step 0 only the non-degraded image is considered, while at time step T only the degraded image is considered. To implement the degradation function $d(\cdot)$, we use part of the BSRGAN [44] framework that creates a random sequence of image mappings that imitate real-life degradations.

Diffusion models are commonly trained on the full range of time steps $[1, T]$ and learn to generate images from pure noise. However, such a setting is not relevant in the context of quality assessments, as the generated (denoised) images have to exhibit a sufficient correspondence with the input samples x . The easiest solution to this issue is to limit the number of time steps, on which the model is trained $t \in [1, T']$, where $T' < T$, and, in turn, ensure that the noisy image is properly conditioned on the input x . The extended training procedure then minimizes (5) until convergence.

3.4. Generating Noisy and Reconstructed Images

To estimate the quality of a given face image x , DiffIQA makes use of the forward and backward diffusion processes of the trained DDPM. Because head pose is an important factor of face quality, which the underlying DDPM can not

Table 1. **Summary of the characteristics of the experimental datasets.** We evaluate DiffFIQA across seven diverse datasets with different quality factors and of different size.

Dataset	#Images	#IDs	#Comparisons		Main Quality Factors ^{†‡}			
			Mated	Non-mated	Pose	O-E	B-R-N	Sc
Adience [11]	19,370	2,284	20,000	20,000	M	M	L	M
CALFW [47]	12,174	4,025	3,000	3,000	M	M	L	M
CFP-FP [37]	7,000	500	3,500	3,500	H	L	L	M
CPLFW [46]	11,652	3,930	3,000	3,000	H	L	M	M
IJB-C [29]	23,124 ^{††}	3,531	19,557	15,638,932	H	H	H	Lr
LFW [19]	13,233	5,749	3,000	3,000	L	L	L	M
XQFW [26]	13,233	5,749	3,000	3,000	L	L	H	M

[†]O-E - Occlusion, Expression; B-R-N - Blur, Resolution, Noise; Sc - Scale.

[‡]L - Low; M - Medium; H - High; Lr - Large; Values estimated subjectively by the authors.

^{††} number of templates, each containing several images

explicitly account for, we extend our methodology, by first constructing a horizontally flipped image x^f that we utilize alongside the original image x in the quality-score calculation step, similarly to [3]. The main intuition behind this approach is to exploit the symmetry of human faces, where large deviations from frontal pose induce large disparities between the embeddings of the original and flipped images that can be quantified during quality estimation. Thus, for a pair of input face images (x, x^f) and a given time step t , we construct a pair of noisy (x_t, x_t^f) and restored images (\hat{x}, \hat{x}^f) and use the generated data for quality estimation.

3.5. Quality-Score Calculation

DiffFIQA relies on the assumption that the embeddings of lower-quality images are more sensitive to image perturbations introduced by the forward and backward diffusion processes than higher-quality images. To quantify this sensitivity, we calculate the average cosine similarity between the embedding of the input image x and all generated noisy and restored counterparts. Additionally, since diffusion models rely on the (random) sampling from a normal distribution, we repeat the whole process n times and average the results, i.e.,

$$q_x = \frac{1}{n|\mathcal{E}|} \sum_{i=1}^n \sum_{e_y \in \mathcal{E}} \frac{e_x^T e_y}{\|e_x\| \cdot \|e_y\|}, \quad (7)$$

where \mathcal{E} is a set of generated image embeddings, i.e., $\mathcal{E} = \{e_{x_t}, e_{\hat{x}}, e_{x_t^f}, e_{\hat{x}^f}\}$, computed with the FR model M as $e_z = M(x_z)$. In the above equation, the operator $\|\cdot\|$ denotes the set cardinality and $\|\cdot\|$ the L_2 norm.

3.6. Model Distillation

One of the main shortcomings of DiffFIQA (and diffusion models in general) is the high computational complexity compared to other types of FIQA techniques. This complexity stems from the iterative nature of the backward diffusion process, which requires a large number of forward passes through the generative network. Since our approach repeats this process n -times, this only exacerbates the problem and adversely affects the applicability of the technique in real-world applications. To address this problem, we

distill the knowledge encoded by DiffFIQA into a regression model. Specifically, we select a pretrained CosFace FR model augmented with a (quality) regression head and fine-tune it on roughly two million quality labels extracted from the VGGFace2 [6] dataset using the proposed DiffFIQA technique. Here, the labels are normalized to $[0, 1]$ and then split into train and validation sets for the training procedure. We refer to the distilled CosFace model as DiffFIQA(R) hereafter, and evaluate it together with the original DiffFIQA technique in the following sections.

4. Experiments and Results

4.1. Experimental Setup

Experimental setting. We analyze the performance of DiffFIQA in comparison to 10 state-of-the-art FIQA methods: (i) the **analytical** FaceQAN [3], SER-FIQ [39], and FaceQgen [15] models, (ii) the **regression-based** FaceQnet [17], SDD-FIQA [32], PCNet [43], and LightQnet [7] techniques, and (iii) the **model-based** MagFace [30], PFE [38], and CR-FIQA [5] methods. We test all methods on 7 commonly used benchmarks with different quality characteristics, as summarized in Table 1, i.e.: Adience [11], Cross-Age Labeled Faces in the Wild (CALFW) [47], Celebrities in Frontal-Profile in the Wild (CFP-FP) [37], Cross-Pose Labeled Faces in the Wild (CPLFW) [46], large-scale IARPA Janus Benchmark C (IJB-C) [29], Labeled Faces in the Wild (LFW) [19] and the Cross-Quality Labeled Faces in the Wild (XQFW) [26]. Because the performance of FIQA techniques is dependent on the FR model used, we investigate how well the techniques generalize over 4 state-of-the-art models, i.e.: AdaFace¹ [24], ArcFace² [10], CosFace² [42], and CurricularFace³ [20] - all named after their training losses. All FR models use a ResNet100 backbone, and are trained on the WebFace12M¹, MS1MV3², Glint360k², and CASIA-WebFace³ datasets.

Evaluation methodology. Following standard evaluation methodology [3, 5, 39] and taking recent insights into account [21, 36], we evaluate the performance of DiffFIQA using *non-interpolated* Error-versus-Discard Characteristic (EDC) curves (often also referred to as Error-versus-Reject Characteristic or ERC curves in the literature) and the consequent pAUC (partial Area Under the Curve) values. The EDC curves measure the False Non-Match Rate (FNMR), given a predefined False Match Rate (FMR) (10^{-3} in our case), with increasing low-quality image discard (reject) rates. In other words, EDC curves measure how the performance of a given FR model improves when some percentage of the lowest quality images is discarded. Since rejecting a large percentage of all samples is not feasible/practical in real-world application scenarios, we are typically most

¹<https://github.com/mk-minchul/AdaFace>

²<https://github.com/deepinsight/insightface>

³<https://github.com/HuangYG123/CurricularFace>

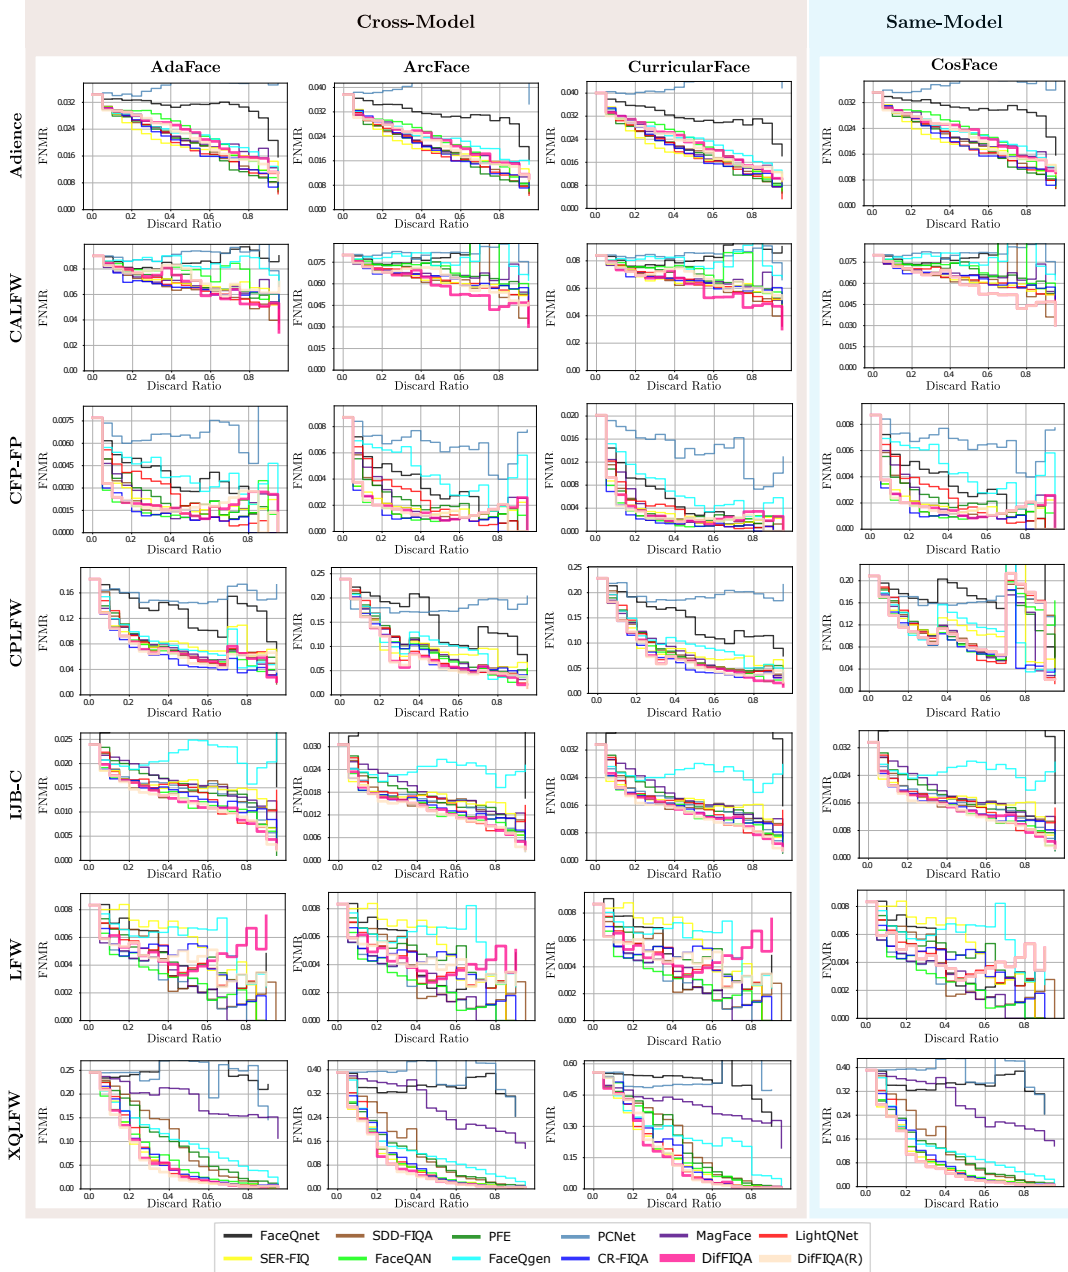


Figure 4. **Comparison to the state-of-the-art with (non-interpolated) EDC curves.** Results are presented for 7 datasets, 4 FR models and with 10 recent FIQA competitors. Observe how the distilled model performs comparably to the non-distilled one, especially at low discard rates. DiffFIQA and DiffFIQA(R) most convincingly outperform the competitors on the most challenging IJB-C and XQFW datasets.

interested in the performance at the lower discard rates. For this reason we report the pAUC values, where only the results up to a predetermined drop rate threshold are considered. Furthermore, for easier interpretation and comparison of scores over different dataset, we normalize the calculated pAUC values using the FNMR at 0% discard rate, with lower pAUC values indicating better performance.

Implementation Details. During training of the DDPM, the maximum number of forward steps is set to $T = 1000$,

yet the underlying model is trained only using up to $T' = 100$ forward diffusion steps. The value of T' does not define the number of time steps t taken at inference time, it only sets the possible upper bound. This process ensures that images produced by the forward process are only partially noisy, so the backward process is properly conditioned on the input image and learns to restore it during training. To account for the randomness introduced by the forward process, we repeat the diffusion process $n = 10$ times and average the results over all iterations, when computing the

final quality score. The utilized UNet model (D) consists of four downsampling and upsampling modules, each decreasing (increasing) the dimensions of the representations by a factor of two. Training is done using the Adam optimizer, with a learning rate of $8.0e^{-5}$ in combination with an Exponential Moving Average (EMA) model, with a decay rate of 0.995. The presented hyperparameters were determined through preliminary experiments on hold-out data to ensure a reasonable trade-off between training speed and reproducible performance. All experiments were conducted on a desktop PC with an Intel i9-10900KF CPU, 64 GB of RAM and an Nvidia 3090 GPU.

4.2. Comparison with the State-of-the-Art

In this section, we compare DiffFIQA and the distilled version, DiffFIQA(R), with ten state-of-the-art competitors and analyze: (i) the *performance characteristics* of the tested techniques, and (ii) their *runtime complexity*.

Performance analysis. In Figure 4, we show the (non-interpolated) EDC curves for all tested FR models and datasets, and report the corresponding pAUC scores in Table 2. Following the suggestions in [21, 36], we chose a discard rate of 0.3, when calculating the pAUC values, but also report additional results in the supplementary material. The experiments are split into the *Cross-* and *Same-Model* scenarios. In the *Cross-Model* scenario the model used to produce the quality score and the model used for evaluation differ, while in the *Same-Model* scenario the used models are the same. We observe that the proposed diffusion-based FIQA techniques result in highly competitive performance across all datasets and FR models. The distilled DiffFIQA(R) model, for example, leads to the lowest average pAUC score with the ArcFace and CosFace FR models, and is the runner-up with the AdaFace and CurricularFace models with pAUC scores comparable to the top performer CR-FIQA⁴. Several interesting findings can be made from the reported results, e.g.: (i) While the performance of DiffFIQA and DiffFIQA(R) is in general close, the distilled version has a slight edge over the original, which suggests that the distillation process infuses some additional information into the FIQA procedure through the FR-based regression model; (ii) The proposed FIQA models are particularly competitive on the difficult large-scale IJB-C dataset, where the DiffFIQA(R) approach consistently outperforms all competing baseline models. A similar observation can also be made for the XQFW dataset, where the diffusion-based models are again the top performers, which speaks of the effectiveness of diffusion-based quality estimation.

Runtime complexity. In Table 3, we compare the runtime complexity of the evaluated FIQA techniques (in ms). To ensure a fair comparison, we utilize (i) the same experimental hardware for all methods (described in Section 4.1), (ii)

Table 2. **Comparison to the state-of-the-art.** The table reports pAUC scores at a discard rate of 0.3 and a FMR of 10^{-3} . Average results across all datasets are marked $\overline{\text{pAUC}}$. The best result for each dataset is shown in bold, the overall best result is colored green, the second best blue and the third best red.

AdaFace - pAUC@FMR= 10^{-3} (↓)								
FIQA model	Adience	CALFW	CFP-PP	CPLFW	IJB-C	LFW	XQFW	pAUC
FaceNet [17]	0.963	0.938	0.717	0.887	1.256	0.884	0.977	0.946
SDD-FIQA [32]	0.839	0.871	0.500	0.688	0.782	0.825	0.842	0.764
PFE [38]	0.833	0.890	0.566	0.681	0.868	0.771	0.798	0.772
PCNet [43]	1.005	0.979	0.862	0.898	0.788	0.661	0.987	0.883
MagFace [30]	0.860	0.866	0.524	0.664	0.883	0.666	0.913	0.768
LightQNet [7]	0.847	0.894	0.641	0.684	0.797	0.777	0.704	0.763
SER-FIQ [†] [40]	0.807	0.892	0.475	0.626	0.762	0.935	n/a	0.749
FaceQAN [3]	0.890	0.919	0.383	0.619	0.756	0.656	0.654	0.697
CR-FIQA [5]	0.844	0.851	0.391	0.588	0.750	0.707	0.684	0.688
FaceQgen [15]	0.858	0.970	0.718	0.694	0.853	0.834	0.736	0.809
DiffFIQA (ours)	0.864	0.900	0.416	0.608	0.761	0.719	0.627	0.699
DiffFIQA(R) (ours)	0.865	0.895	0.412	0.601	0.731	0.708	0.610	0.689
ArcFace - pAUC@FMR= 10^{-3} (↓)								
FIQA model	Adience	CALFW	CFP-PP	CPLFW	IJB-C	LFW	XQFW	pAUC
FaceNet [17]	0.943	0.955	0.693	0.878	1.224	0.884	0.899	0.925
SDD-FIQA [32]	0.783	0.901	0.491	0.734	0.720	0.808	0.774	0.744
PFE [38]	0.774	0.932	0.524	0.738	0.783	0.779	0.641	0.739
PCNet [43]	1.022	1.006	0.868	0.783	0.706	0.623	1.004	0.859
MagFace [30]	0.812	0.902	0.549	0.717	0.824	0.635	0.943	0.769
LightQNet [7]	0.789	0.913	0.612	0.752	0.721	0.745	0.621	0.736
SER-FIQ [†] [39]	0.767	0.903	0.416	0.656	0.671	0.935	n/a	0.724
FaceQAN [3]	0.824	0.941	0.373	0.677	0.673	0.624	0.581	0.670
CR-FIQA [5]	0.808	0.891	0.358	0.689	0.664	0.675	0.642	0.675
FaceQgen [15]	0.817	0.985	0.784	0.701	0.785	0.802	0.653	0.789
DiffFIQA (ours)	0.805	0.900	0.399	0.647	0.675	0.695	0.546	0.667
DiffFIQA(R) (ours)	0.801	0.898	0.389	0.646	0.655	0.708	0.554	0.665
CurricularFace - pAUC@FMR= 10^{-3} (↓)								
FIQA model	Adience	CALFW	CFP-PP	CPLFW	IJB-C	LFW	XQFW	pAUC
FaceNet [17]	0.921	0.947	0.601	0.867	1.248	0.908	0.984	0.925
SDD-FIQA [32]	0.776	0.900	0.409	0.696	0.721	0.821	0.817	0.734
PFE [38]	0.759	0.923	0.415	0.691	0.784	0.785	0.835	0.742
PCNet [43]	1.004	0.996	0.887	0.899	0.710	0.656	0.938	0.870
MagFace [30]	0.793	0.892	0.477	0.689	0.821	0.661	0.862	0.742
LightQNet [7]	0.769	0.910	0.462	0.704	0.713	0.767	0.739	0.723
SER-FIQ [†] [39]	0.750	0.883	0.389	0.625	0.661	0.942	n/a	0.708
FaceQAN [3]	0.811	0.931	0.343	0.637	0.669	0.644	0.835	0.696
CR-FIQA [5]	0.797	0.877	0.318	0.615	0.664	0.693	0.789	0.679
FaceQgen [15]	0.815	0.974	0.662	0.698	0.783	0.845	0.750	0.790
DiffFIQA (ours)	0.806	0.884	0.384	0.624	0.672	0.711	0.736	0.688
DiffFIQA(R) (ours)	0.788	0.892	0.358	0.622	0.644	0.724	0.768	0.685
CosFace - pAUC@FMR= 10^{-3} (↓)								
FIQA model	Adience	CALFW	CFP-PP	CPLFW	IJB-C	LFW	XQFW	pAUC
FaceNet [17]	0.952	0.955	0.693	0.879	1.248	0.884	0.899	0.930
SDD-FIQA [32]	0.825	0.901	0.491	0.735	0.721	0.808	0.774	0.751
PFE [38]	0.813	0.932	0.524	0.748	0.784	0.779	0.641	0.746
PCNet [43]	1.009	1.006	0.868	0.835	0.710	0.623	1.004	0.865
MagFace [30]	0.852	0.902	0.549	0.724	0.821	0.635	0.943	0.775
LightQNet [7]	0.835	0.913	0.612	0.753	0.713	0.745	0.621	0.742
SER-FIQ [†] [39]	0.793	0.903	0.416	0.711	0.661	0.935	n/a	0.736
FaceQAN [3]	0.871	0.941	0.373	0.667	0.675	0.624	0.581	0.676
CR-FIQA [5]	0.835	0.891	0.358	0.681	0.664	0.675	0.642	0.678
FaceQgen [15]	0.847	0.985	0.784	0.702	0.783	0.802	0.653	0.794
DiffFIQA (ours)	0.841	0.900	0.399	0.669	0.672	0.695	0.546	0.675
DiffFIQA(R) (ours)	0.838	0.900	0.389	0.669	0.644	0.695	0.546	0.669

[†]SER-FIQ was used to create XQFW, so the results here are not reported for a fair comparison.

use the official code, released by the authors for all techniques, and (iii) compute average runtimes and standard deviations over the entire XQFW dataset. As can be seen, the original approach, DiffFIQA, despite being highly competitive in terms of performance, is among the most computationally demanding due to the use of the complex diffusion processes. With around 1s on average per image, the runtime complexity of the model is even significantly higher than that of the FaceQAN or SER-FIQ techniques that require multiple passes through their networks to estimate quality and which are already among the slower FIQA models. However, the distillation process, allows to reduce the runtime by roughly three orders of magnitude (or by 99,9%), making the distilled DiffFIQA(R) comparable to the faster models evaluated in this experiment.

⁴In the supplementary material we show that with a discard rate of 0.2, DiffFIQA(R) is the top performer with 3 of the 4 FR models.

Table 3. **Runtime complexity.** The reported results (in ms) were computed over the XQLFW dataset and the same experimental hardware. Note how the distillation process leads to a speed-up of more than three orders of magnitude from DiffFIQA to DiffFIQA(R).

FIQA model	FaceNet [17]	SDD-FIQA [32]	PFE [38]	PCNet [43]	MagFace [30]	LightQNet [7]	SER-FIQA [39]	FaceQAN [3]	CR-FIQA [5]	FaceQgen [15]	Ours	
											DiffFIQA	DiffFIQA(R)
Runtime ($\mu \pm \sigma$)	42.11 \pm 2.14	0.62 \pm 0.36	42.69 \pm 12.26	17.06 \pm 0.34	1.08 \pm 0.36	18.54 \pm 18.68	112.93 \pm 33.81	334.13 \pm 118.79	0.15 \pm 0.37	42.11 \pm 2.05	1074.62 \pm 11.45	1.24 \pm 0.36

Table 4. **Results of the ablation study.** The results are reported in terms of pAUC (\downarrow) at a FMR of 10^{-3} and a discard rate of 0.3.

Model variant	LFW	CPLFW	CALFW	XQLFW	pAUC
(A1): w/o Image Flipping	0.702	0.727	0.888	0.535	0.713
(A2): w/o Forward Pass	0.730	0.684	0.897	0.531	0.710
(A3): DiffFIQA ($t = 20$)	0.657	0.694	0.945	0.628	0.731
DiffFIQA (complete)	0.695	0.669	0.900	0.546	0.702




Figure 5. **Quality-score distributions.** DiffFIQA and DiffFIQA(R) produce very consistent distributions over all seven test datasets.

4.3. Ablation Study

We perform several ablation studies to explore the impact of the main components of DiffFIQA. Specifically, we are interested in: (A1) the impact of the flipping procedure, utilized to capture pose-related quality factors, (A2) the contribution of the forward pass (i.e., the noising step of the diffusion), and (A3) the impact of the number of forward diffusion steps t , where a larger number corresponds to higher amounts of noise in the image x_t produced by the forward process. Because the ablations are only relevant for the (non-distilled) approach, we experiment solely with the DiffFIQA technique and report results using the CosFace FR model and four datasets that feature a broad range of quality characteristics, i.e., LFW, CPLFW, CALFW and XQLFW.

From the results in Table 4, we observe that the exclusion of the flipping operation significantly degrades performance on the cross-pose (CPLFW) dataset, while contributing to minor improvements on CALFW and XQLFW. However, given that pose is considered one of the main factors still adversely affecting modern FR models, the flipping operation still helps with the performance across all the test datasets (see average A1 results). When removing the forward pass (in A2), we again see considerable performance drops on LFW and CPLFW, leading to lower average pAUC scores. This suggest that both (forward and backward) processes are important for good results across different datasets. Finally, we see that lesser amounts of noise and stronger conditioning on the input images leads to better results as evidenced by the A3 results with our model with 20 timesteps, instead of the 5 utilized in the complete DiffFIQA approach.




Figure 6. **Illustration of the quality scores produced by the proposed FIQA techniques.** The scores on the top shows results for DiffFIQA and the scores at the bottom for DiffFIQA(R). While the concrete scores differ, both models generate similar rankings.

4.4. Qualitative Evaluation

While the proposed DiffFIQA approach has a sound theoretical basis that links the forward and backward diffusion processes to face-image quality, the distilled variant abstracts this relation away and approaches the FIQA task from a pure learning perspective. To get better insight into the characteristics of both models, we investigate in this section their behavior in a qualitative manner.

Quality-score distributions. In Figure 5, we compare the quality-score distributions, generated by DiffFIQA and DiffFIQA(R) on all seven test datasets. As can be seen, the two models produce very similar distributions, with a slight preference of DiffFIQA(R) towards higher quality scores.

Visual analysis. In Figure 6, we show sample images from XQLFW and the corresponding quality scores, generated by DiffFIQA and DiffFIQA(R). Note that both approaches produce a similar ranking but differ in the concrete quality score assigned to a given image. Interestingly, some blurry images with low (human-perceived) visual quality receive relatively high quality scores, as they feature frontal faces that may still be useful for recognition purposes. Additional qualitative results that illustrate the capabilities of the DiffFIQA model are also shown on the right part of Figure 1.

5. Conclusion

We have presented a novel FIQA approach, called DiffFIQA, that uses denoising diffusion probabilistic models as the basis for quality estimation. Through comprehensive experiments on multiple datasets we showed that the proposed model yields highly competitive results, when benchmarked against current state-of-the-art techniques and that the runtime performance can be reduced significantly if the model is distilled into a quality predictor through a regression-based procedure. Going forward, we plan to investigate extensions to our model, including transformer-based UNet alternatives and latent diffusion processes.

Acknowledgments. ARRS grants P2-0250 and J2-2501.

References

- [1] S. Anwarul and S. Dahiya. A Comprehensive Review on Face Recognition Methods and Factors Affecting Facial Recognition Accuracy. *Proceedings of International Conference on Recent Innovations in Computing (ICRIC)*, pages 495–514, 2020.
- [2] Ž. Babnik, D. Naser, and V. Štruc. Optimization-Based Improvement of Face Image Quality Assessment Techniques. In *Proceedings of the International Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2023.
- [3] Ž. Babnik, P. Peer, and V. Štruc. FaceQAN: Face Image Quality Assessment through Adversarial Noise Exploration. In *Proceedings of the IAPR International Conference on Pattern Recognition (ICPR)*, pages 748–754, 2022.
- [4] L. Best-Rowden and A. K. Jain. Learning Face Image Quality from Human Assessments. *Transactions on Information Forensics and Security (TIFS)*, 13(12):3064–3077, 2018.
- [5] F. Boutros, M. Fang, M. Klemm, B. Fu, and N. Damer. CR-FIQA: Face Image Quality Assessment by Learning Sample Relative Classifiability. In *Proceedings of the CVF/IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [6] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. Vggface2: A Dataset for Recognising Faces Across Pose and Age. In *Proceedings of the IEEE International Conference on Automatic Face & Gesture Recognition (FG)*, pages 67–74, 2018.
- [7] K. Chen, T. Yi, and Q. Lv. LightQNet: Lightweight Deep Face Quality Assessment for Risk-Controlled Face Recognition. *Signal Processing Letters*, 28:1878–1882, 2021.
- [8] K. Chen, T. Yi, and Q. Lv. Fast and Reliable Probabilistic Face Embeddings Based on Constrained Data Uncertainty Estimation. *Image and Vision Computing (IVC)*, page 104429, 2022.
- [9] F.-A. Croitoru, V. Hondru, R. T. Ionescu, and M. Shah. Diffusion Models in Vision: A survey. *Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2023.
- [10] J. Deng, J. Guo, N. Xue, and S. Zafeiriou. Arcface: Additive Angular Margin Loss for Deep Face Recognition. In *Proceedings of the CVF/IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4690–4699, 2019.
- [11] E. Eidinger, R. Enbar, and T. Hassner. Age and Gender Estimation of Unfiltered Faces. *Transactions on Information Forensics and Security (TIFS)*, 9(12):2170–2179, 2014.
- [12] X. Gao, S. Z. Li, R. Liu, and P. Zhang. Standardization of Face Image Sample Quality. In *Proceedings of the IAPR International Conference on Biometrics (ICB)*, pages 242–251. Springer, 2007.
- [13] K. Grm, V. Štruc, A. Artiges, M. Caron, and H. K. Ekenel. Strengths and weaknesses of deep learning models for face recognition against image degradations. *IET Biometrics*, 7(1):81–89, 2018.
- [14] O. Henniger, B. Fu, and C. Chen. On the Assessment of Face Image Quality Based on Handcrafted Features. In *Proceedings of the IEEE International Conference of the Biometrics Special Interest Group (BIOSIG)*, pages 1–5, 2020.
- [15] J. Hernandez-Ortega, J. Fierrez, I. Serna, and A. Morales. FaceQgen: Semi-Supervised Deep Learning for Face Image Quality Assessment. In *Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–8, 2021.
- [16] J. Hernandez-Ortega, J. Galbally, J. Fierrez, and L. Beslay. Biometric Quality: Review and Application to Face Recognition with FaceQnet. *arXiv preprint arXiv:2006.03298*, 2020.
- [17] J. Hernandez-Ortega, J. Galbally, J. Fierrez, R. Haraksim, and L. Beslay. FaceQnet: Quality Assessment for Face Recognition Based on Deep Learning. In *Proceedings of the IAPR International Conference on Biometrics (ICB)*, pages 1–8, 2019.
- [18] J. Ho, A. Jain, and P. Abbeel. Denoising Diffusion Probabilistic Models. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:6840–6851, 2020.
- [19] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007.
- [20] 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 *Proceedings of the CVF/IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5901–5910, 2020.
- [21] ISO/IEC DIS 29794-1, Biometric Sample Quality. Standard, International Organization for Standardization (ISO), 2022.
- [22] M. Ivanovska and V. Štruc. Face Morphing Attack Detection with Denoising Diffusion Probabilistic Models. In *Proceedings of the International Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2023.
- [23] H.-I. Kim, S. H. Lee, and Y. M. Ro. Investigating Cascaded Face Quality Assessment for Practical Face Recognition System. In *Proceedings of the IEEE International Symposium on Multimedia (ISM)*, pages 399–400, 2014.
- [24] M. Kim, A. K. Jain, and X. Liu. AdaFace: Quality Adaptive Margin for Face Recognition. In *Proceedings of the CVF/IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18750–18759, 2022.
- [25] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [26] M. Knoche, S. Hormann, and G. Rigoll. Cross-Quality LFW: A Database for Analyzing Cross-Resolution Image Face Recognition in Unconstrained Environments. In *Proceedings of the IEEE International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–5, 2021.
- [27] S. Li, J. Xu, X. Xu, P. Shen, S. Li, and B. Hooi. Spherical Confidence Learning for Face Recognition. In *Proceedings of the CVF/IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15629–15637, 2021.
- [28] Z. Lijun, S. Xiaohu, Y. Fei, D. Pingling, Z. Xiangdong, and S. Yu. Multi-Branch Face Quality Assessment for Face Recognition. In *Proceedings of the IEEE International Conference on Communication Technology (ICCT)*, pages 1659–1664, 2019.

- [29] B. Maze, J. Adams, J. A. Duncan, N. Kalka, T. Miller, C. Otto, A. K. Jain, W. T. Niggel, J. Anderson, J. Cheney, et al. IARPA Janus Benchmark-C: Face Dataset and Protocol. In *Proceedings of the International Conference on Biometrics (ICB)*, pages 158–165, 2018.
- [30] Q. Meng, S. Zhao, Z. Huang, and F. Zhou. MagFace: A Universal Representation for Face Recognition and Quality Assessment. In *Proceedings of the CVF/IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14225–14234, 2021.
- [31] K. Nasrollahi and T. B. Moeslund. Extracting a Good Quality Frontal Face Image From a Low-Resolution Video Sequence. *Transactions on Circuits and Systems for Video Technology (TCSVT)*, 21(10):1353–1362, 2011.
- [32] F.-Z. Ou, X. Chen, R. Zhang, Y. Huang, S. Li, J. Li, Y. Li, L. Cao, and Y.-G. Wang. SDD-FIQA: Unsupervised Face Image Quality Assessment with Similarity Distribution Distance. In *Proceedings of the CVF/IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7670–7679, 2021.
- [33] R. Raghavendra, K. B. Raja, B. Yang, and C. Busch. Automatic Face Quality Assessment From Video Using Gray Level Co-Occurrence Matrix: An Empirical Study on Automatic Border Control System. In *Proceedings of the IAPR International Conference on Pattern Recognition (ICPR)*, pages 438–443, 2014.
- [34] O. Ronneberger, P. Fischer, and T. Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, pages 234–241, 2015.
- [35] T. Schlett, C. Rathgeb, O. Henniger, J. Galbally, J. Fierrez, and C. Busch. Face Image Quality Assessment: A Literature Survey. *ACM Computing Surveys (CSUR)*, 54(10s):1–49, 2022.
- [36] T. Schlett, C. Rathgeb, J. Tapia, and C. Busch. Considerations on the Evaluation of Biometric Quality Assessment Algorithms. *arXiv preprint arXiv:2303.13294*, 2023.
- [37] S. Sengupta, J. C. Cheng, C. D. Castillo, V. M. Patel, R. Chellappa, and D. W. Jacobs. Frontal to Profile Face Verification in the Wild. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2016.
- [38] Y. Shi and A. K. Jain. Probabilistic Face Embeddings. In *Proceedings of the CVF/IEEE International Conference on Computer Vision (CVPR)*, pages 6902–6911, 2019.
- [39] P. Terhorst, J. N. Kolf, N. Damer, F. Kirchbuchner, and A. Kuijper. SER-FIQ: Unsupervised Estimation of Face Image Quality Based on Stochastic Embedding Robustness. In *Proceedings of the CVF/IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5651–5660, 2020.
- [40] P. Terhorst, J. N. Kolf, N. Damer, F. Kirchbuchner, and A. Kuijper. SER-FIQ: Unsupervised Estimation of Face Image Quality Based on Stochastic Embedding Robustness. In *Proceedings of the CVF/IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5651–5660, 2020.
- [41] P. Terhörst, J. N. Kolf, M. Huber, F. Kirchbuchner, N. Damer, A. M. Moreno, J. Fierrez, and A. Kuijper. A Comprehensive Study on Face Recognition Biases Beyond Demographics. *Transactions on Technology and Society (TTS)*, 3(1):16–30, 2021.
- [42] 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 *Proceedings of the CVF/IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5265–5274, 2018.
- [43] W. Xie, J. Byrne, and A. Zisserman. Inducing Predictive Uncertainty Estimation for Face Verification. In *British Machine Vision Conference (BMVC)*, 2020.
- [44] K. Zhang, J. Liang, L. Van Gool, and R. Timofte. Designing a Practical Degradation Model for Deep Blind Image Super-Resolution. In *Proceedings of the CVF/IEEE International Conference on Computer Vision (ICCV)*, pages 4791–4800, 2021.
- [45] Y. Zhao and W. Deng. Dual Gaussian Modeling for Deep Face Embeddings. *Pattern Recognition Letters*, 161:74–81, 2022.
- [46] T. Zheng and W. Deng. Cross-Pose LFW: A Database for Studying Cross-Pose Face Recognition in Unconstrained Environments. Technical Report 18-01, Beijing University of Posts and Telecommunications, February 2018.
- [47] T. Zheng, W. Deng, and J. Hu. Cross-Age LFW: A Database for Studying Cross-Age Face Recognition in Unconstrained Environments. *CoRR*, abs/1708.08197, 2017.