

Optimization-Based Improvement of Face Image Quality Assessment Techniques

Žiga Babnik*, Naser Damer†, Vitomir Štruc*

*University of Ljubljana, Faculty of Electrical Engineering Tržaška cesta 25, 1000 Ljubljana, Slovenia

†Fraunhofer IGD, Darmstadt, Germany

Abstract—Contemporary face recognition (FR) models achieve near-ideal recognition performance in constrained settings, yet do not fully translate the performance to unconstrained (real-world) scenarios. To help improve the performance and stability of FR systems in such unconstrained settings, face image quality assessment (FIQA) techniques try to infer sample-quality information from the input face images that can aid with the recognition process. While existing FIQA techniques are able to efficiently capture the differences between high and low quality images, they typically cannot fully distinguish between images of similar quality, leading to lower performance in many scenarios. To address this issue, we present in this paper a supervised quality-label optimization approach, aimed at improving the performance of existing FIQA techniques. The developed optimization procedure infuses additional information (computed with a selected FR model) into the initial quality scores generated with a given FIQA technique to produce better estimates of the “actual” image quality. We evaluate the proposed approach in comprehensive experiments with six state-of-the-art FIQA approaches (CR-FIQA, FaceQAN, SER-FIQ, PCNet, MagFace, SER-FIQ) on five commonly used benchmarks (LFW, CFP-FP, CPLFW, CALFW, XQFW) using three targeted FR models (ArcFace, ElasticFace, CurricularFace) with highly encouraging results.

Index Terms—Biometrics, Face recognition, Face image quality assessment, Optimization, Transfer learning

I. INTRODUCTION

Modern face recognition (FR) systems achieve excellent results even with large-scale recognition problems, as long as the appearance variability of the facial images is reasonably constrained. However, the performance in constrained scenarios does not always translate to real-world scenarios where out-of-distribution data, often of poor quality, still presents a challenge for the majority of existing FR models [1], [2].

Face image quality assessment (FIQA) techniques aim to assist FR models in such challenging scenarios by providing additional information on the quality of facial images. This quality information can then be used to either reject low-quality samples that typically lead to false match errors or design robust quality-aware face recognition techniques. Thus, different from general purpose image quality assessment (IQA) methods [3]–[5] that commonly measure the perceived visual quality of images by examining explicit image characteristics, such as sharpness, lighting conditions and resolution, FIQA techniques typically try to capture the utility (or fitness) of the given face image for the recognition task [6]. In other words, they measure the usefulness of the sample for face recognition.

Several groups of FIQA techniques that differ slightly in their approach have been proposed so far in the literature [7]. The majority of recent techniques learns quality-estimation

networks using (reference) quality information inferred from a large database of face images [8]–[11]. Another notable group of FIQA techniques estimates quality based only on the information present in the input image and the characteristics of the targeted FR system [12], [13]. More recently, approaches have also appeared that incorporate quality estimation directly into the FR process [14], [15], paving the way towards quality-aware face recognition.

While most of the existing FIQA techniques perform well enough to distinguish between high-quality and low-quality facial images, correctly ranking face images of similar quality remains an open problem. The correct (optimal) ordering does not depend solely on the input face images, but also on the targeted FR model. Each model may, in a sense, *perceive* the quality of individual samples differently due to different model-specific biases introduced by the learning process and the data used for training [16], [17]. This observation also suggests that FIQA techniques, that are not FR model specific, can not determine the correct order for all possible FR models. For this reason, we propose in this paper a novel optimization approach, that attempts to improve the predictive power of any given FIQA approach by incorporating quality information obtained by a particular FR model into the quality scores generated by the selected FIQA approach. Thus, the main contributions of this papers are:

- A novel optimization approach that incorporates model-specific quality information into the quality scores produced by existing FIQA techniques with the goal of improving FIQA performance.
- An in-depth evaluation of the proposed optimization approach over six FIQA techniques, five dataset, three recognition models and in two experimental settings that demonstrates a new state-of-the-art for the FIQA task.

II. RELATED WORK

In this section, we briefly review previous FIQA research that can be broadly categorized into three groups: (i) analytical, (ii) regression and (iii) model-based techniques. More in-depth information on face quality assessment can be found in the comprehensive survey paper by Schlett *et al.* [7].

Analytical FIQA techniques are mostly unsupervised and rely solely on the information that can be extracted directly from the given input sample. Techniques from this group typically focus on the visual quality of the facial images and, as a result, often exhibit limited performance. The method proposed by Gao *et al.* [18], for example, attempts to extract quality information based on facial symmetry estimation only. Zhang *et al.* [19] try to quantify quality based on image




Fig. 1. **Overview of the proposed method that consists of: Label Optimization and Transfer Learning.** The *label-optimization* step incorporates information extracted from mated image pairs into quality scores precomputed with an existing FIQA technique. The *transfer-learning* step is then used to train a FR model, extended with a regression head, on the optimized quality-scores. The learned regressor is finally utilized for quality estimation.

illumination information, while Lijun *et al.* [20] combine multiple cues, such as occlusions, blur and pose for the quality-estimation task. Different from these methods, two analytical FIQA techniques have been proposed recently that in addition to the characteristics of the input image also consider the targeted FR system during the quality estimation task. The first, SER-FIQ by Terh rst *et al.* [12], uses the properties of dropout layers to quantify quality, while FaceQAN, by Babnik *et al.* [13], exploits adversarial examples for quality assessment. Both methods were shown to yield state-of-the-art performance for various FR models and different benchmarks.

Regression-based FIQA techniques are the most numerous and usually learn a quality estimation (regression) model to predict quality scores based on some pseudo (ground-truth) quality labels. FaceQNet [8], for example, trains a ResNet50 model using labels obtained by embedding comparisons with the highest quality image of each subject. Here, the highest quality images are determined using an external quality compliance tool. A similar approach, called PCNet [11], trains a quality-regression network on mated-image pairs, with the goal of predicting the similarity of the image pair. LightQNet [10] builds on the ideas introduced with PCNet, but additionally relies on a so-called Identification Quality (IQ) loss, while SDD-FIQA [9] considers both mated and non-mated similarity scores between a large number of samples to determine the final reference quality for the regression task.

Model-based FIQA techniques are less common and usually try to combine face recognition and quality assessment in a single quality-aware face recognition task. The main goal of these techniques is to simultaneously produce, for a given sample, its embedding and an estimate of the sample’s quality. For example, the approach presented by Shi and Jain [14], estimates a mean and variance vector for each given input sample, where the mean vector represents the embedding, while the variance provides the corresponding uncertainty and can be interpreted as a sample quality estimate. MagFace [15], a similar approach by Meng *et al.*, uses a modified version of the commonly used ArcFace loss, called MagFace loss, which is able to generate quality-aware embeddings, by incorporating quality information into the magnitude of the embedding itself.

The method we propose cannot be clearly assigned to one of the above groups, because it relies on an already existing FIQA approach (from any of the three groups) to generate reference quality scores. In a sense, it distills FIQA knowledge from any existing technique. However, if treated as a black-box, the

proposed FIQA approach can be thought of as a regression-based technique, as it trains a regression model using quality labels extracted from a large database.

III. METHODOLOGY

State-of-the-art FIQA techniques are able to efficiently discriminate between images of distinctly different qualities, yet may not be able to properly distinguish between images of similar quality. To exacerbate this problem, the relative ordering of images of similar quality may additionally depend on the targeted FR model, which not all FIQA techniques take into account. Because face quality assessment aims to quantify the utility of face images for a given FR model, the slight variations in the biases present in modern FR systems may result in different (optimal) quality scores for different FR models. For this reason, we propose in this paper an approach that aims to incorporate FR model-specific quality information into (some initial) quality scores, with the goal of improving the fine-grained performance of existing FIQA techniques. The overall pipeline of the proposed approach, shown in Fig. 1, consists of two main steps: (i) *label optimization* and (ii) *transfer learning*. The label-optimization step aims to incorporate additional quality-related information into the baseline quality labels, precomputed using a selected (existing) FIQA approach. The optimized quality labels are then used in a transfer-learning scheme, that uses a pre-trained FR model, extended with a quality-regression head.

A. Method Overview

Let Q and M denote a given FIQA method and a pre-trained FR model that produce quality scores $q_I = Q(I)$ and embeddings $e_I = M(I)$, respectively, for an arbitrary input face image I , and $\{I_i\}_{i=1}^N$ denote a large facial image database consisting of N distinct images. The goal of our approach is to train a regression-based quality-estimation model $Q^* = H(M(I))$, where H represents a quality-regression head, that outperforms the initial FIQA method Q . The model Q^* is trained on optimized quality labels $\{q_i^*\}_{i=1}^N$ generated by the proposed optimization scheme O . The method relies on information obtained from mated image pairs of the face database $\{I_i\}_{i=1}^N$. Details on the procedure are given below.

B. Initialization

We first extract initial quality scores $q_i = Q(I_i)$ and embeddings $e_i = M(I_i)$ from all images of the given face




Fig. 2. **Overview of Label Optimization.** We present a visualization of the proposed optimization scheme. Based on the embeddings $\{e_{I_i}\}_{i=1}^{N_k}$ we first generate mated image pairs. From the image pairs, we compute the pair similarity distribution \mathcal{X}_s using the cosine similarity of the image embeddings. At the same time, we also construct the quality distribution \mathcal{X}_q from the given quality scores $\{q_{I_i}\}_{i=1}^N$. The mean similarity index $\bar{id}_{I_i}^s$, calculated as the average index of all image pairs from \mathcal{I} , is then used to update the quality index $id^{\mathcal{X}_q}(q_{I_i})$, using the equation presented above.

image database $\{I_i\}_{i=1}^N$ using the selected FIQA method Q and chosen FR model M . This initialization step is conducted once and provides the input data for the label optimization and consequently the transfer learning procedures.

C. Label Optimization

Looking at past research [8]–[11], we observe that quality information is often inferred from mated image comparisons, where the term *mated images* refers to two unique images of the same individual. We, therefore, follow this insight and use such information in our optimization approach as well. By computing the similarity of mated image pairs in the embedding space of the given FR model M , we are also able to include FR-specific quality estimates in the optimization.

Selecting mated image pairs. Large-scale databases contain a significant amount of images for each individual, where many of the images may be nearly identical. Selecting all possible mated pairs, can, therefore, introduce database specific biases into our approach and adversely affect performance.

To avoid such issues, we propose a technique for sampling mated image pairs based on clustering. We use a clustering procedure to find groups of similar images and to identify the most informative (and least redundant) mated image pairs. We cluster the embedding space $\mathcal{E}^k = \{e_i^k\}_{i=1}^{N_k}$ corresponding to images of each individual $k = 1, \dots, K$ present in the database using K-Means, where $N = \sum_k N_k$. The algorithm initializes C cluster centers by randomly sampling the given data points and iteratively corrects them using nearby examples. For each image I_c^k of the k -th individual belonging to cluster $c \in [1, C]$, we randomly select images from all other clusters $c' \neq c$, $c' \in [1, C]$ to form mated pairs $(I_c^k, I_{c'}^k)$. By repeating this process for each image of every individual, we obtain the final mated image pairs for the label-optimization procedure $G = \{(I_i, I_j)^l\}_{l=1}^L$, where $i \neq j$ and $L = N \cdot (C - 1)$.

Optimizing prior quality scores. We aim at optimizing the initial quality labels $\{q_{I_i}\}_{i=1}^N$ using information provided by the average pair similarity sim_{I_i} of each image. In other words, if an image has a low quality score, yet its average pair similarity is high, we want to increase its quality. Conversely, if the opposite is true, we want to decrease it. The design of the optimization procedure is based on the assumption that the initial quality scores already provide a reasonable estimate of the true quality. We, therefore, try to retain the overall quality distribution over the face database. As a result, we

simply rearrange the order of the images in the original quality score distribution generated by the selected FIQA technique Q instead of computing new optimal quality scores that could differ significantly from the initial estimates.

From the list of genuine image pairs G , we first calculate the cosine similarity of all image embedding pairs, i.e.:

$$sim_{cos}(e_{I_i}, e_{I_j}) = \frac{e_{I_i} \cdot e_{I_j}}{\|e_{I_i}\| \cdot \|e_{I_j}\|}, \quad (1)$$

where e_{I_i} and e_{I_j} denote embeddings of images I_i and I_j . We then construct the distribution of the computed similarity scores $\mathcal{X}_s \sim sim_{cos}(e_i, e_j)$, $\forall (I_i, I_j) \in G$, by sorting all the pairs according to their calculated similarity score. From the distribution \mathcal{X}_s we compute for each image I_i its average pair index,

$$\bar{id}_{I_i}^s = \frac{1}{|\mathcal{I}|} \sum_{\mathcal{I}} id^{\mathcal{X}_s}(I_i, I_j), \quad (2)$$

where $id^{\mathcal{X}_s}(\cdot)$ is a function, that for a given pair (I_i, I_j) returns the index of $sim_{cos}(e_i, e_j)$ within the similarity distribution \mathcal{X}_s and \mathcal{I} represents the set of all image pairs (I_i, I_j) , where the quality q_{I_i} is lower than q_{I_j} . The latter follows from the fact that the quality of an image pair is computed as $q(I_i, I_j) = \min(I_i, I_j)$, i.e., it depends only on the image with the lower quality. In addition, we construct a quality score distribution $\mathcal{X}_q \sim \{q_{I_i}\}_{i=1}^N$, by sorting the quality scores of all images within the given database. The average pair indices and the distribution \mathcal{X}_q are then used to compute the optimized quality indices

$$id^{\mathcal{X}_q}(q_{I_i}^*) = id^{\mathcal{X}_q}(q_{I_i}) + \theta \cdot (\bar{id}_{I_i}^s - id^{\mathcal{X}_q}(q_{I_i})), \quad (3)$$

where θ is an open hyperparameter that controls the degree of change for the indices, and $id^{\mathcal{X}_q}(\cdot)$ is a function that returns, for some quality q , its index within the distribution \mathcal{X}_q .

Final steps. To avoid bias from randomly selecting mated pairs, we also repeat the entire process R times, and average the final optimized quality indices, $\bar{id}(q_{I_i}^*) = \frac{1}{R} \sum_{r=1}^R id_r^{\mathcal{X}_q}(q_{I_i}^*)$ for all images. The images are then sorted by the calculated optimized quality indices $\bar{id}(q_{I_i}^*)$ and assigned the quality score according to the output of the sorted list and the original quality score distribution \mathcal{X}_q .

D. Transfer Learning

One of the main goals of FIQA techniques is to improve the stability and performance of FR systems. We propose to use




Fig. 3. **Example VGGFace2 images.** Images of three distinct individuals are shown, illustrating the amount of variability present in the database.

a pre-trained state-of-the-art FR model for quality prediction, as it efficiently extracts identity information from given facial images. Moreover, the embeddings generated by state-of-the-art FR models already contain some information about the quality of the input image. Formally, from an FR model M , we construct a quality regression model $H \circ M$, where H represents a regression head. The regression head H attempts to extract the quality of the input image $q_i = H(e_{I_i})$ from the embedding $e_{I_i} = M(I_i)$ and is learned through an L_1 loss applied over the optimized labels. To improve the transfer-learning process, we normalize the optimized quality scores to the interval of $[0, 1]$.

IV. EXPERIMENTS AND RESULTS

A. Experimental Setup

Training Database. To train the proposed approach, a large-scale database of diverse facial images with rich appearance variability is needed. To this end, we select the VGGFace2 database [21], which contains over 3 million images of more than 9000 individuals. Images in the database vary in terms of facial pose, lighting conditions, image resolution, occlusions, and other similar factors that greatly affect the overall quality and appearance of the facial images, as also illustrated in Fig. 3 for three individuals (in columns) from the database.

Evaluation Setting. We use six state-of-the-art FIQA methods as baselines to evaluate the proposed optimization scheme, i.e., CR-FIQA [22], FaceQAN [13], MagFace [15], SER-FIQ [12], SDD-FIQA [9] and PCNet [11]. The baselines and the learned quality-regression networks are evaluated on five commonly used benchmark databases: XQFW [23], CPLFW [24], CFP-FP [25], CALFW [26] and LFW [27]. As the pre-trained FR model, we use ArcFace [28] with a ResNet100 architecture, trained on the MS1MV3 database using an angular margin loss. For the performance evaluation we consider two different scenarios: (i) the *same-model scenario*, where we use the ArcFace model for both quality-score prediction and generation of the performance indicators, and (ii) the *cross-model scenario* where ArcFace is used for quality assessment, and the CurricularFace [29] and ElasticFace [30] models are utilized to evaluate performance. Both of the test models are based on the ResNet100 architecture, but CurricularFace was trained on MS1MV2, while ElasticFace was trained with CASIA-WebFaces and MS1MV2.

Performance Evaluation. We follow standard FIQA evaluation methodology and use Error-versus-Reject-Characteristic (ERC) curves as the basis for performance reporting [7], [13], [15], [22]. ERC curves measure the False Non-Match Rate

TABLE I
SAME-MODEL PERFORMANCE COMPARISON. COMPARISON OF AUC SCORES BETWEEN BASELINE FIQA METHODS (*Baseline*) AND OUR PROPOSED OPTIMIZATION APPROACH (*Optimized*) USING ARCFACE FOR QUALITY ESTIMATION AND PERFORMANCE EVALUATION. BEST RESULTS ARE HIGHLIGHTED IN **GREEN**.

Methods		AUC@FMR1e-3[$\times 10^{-3}$](\downarrow)				
		LFW	CFP-FP	CPLFW	CALFW	XQFW
CR-FIQA	Baseline	1.7	1.3	86.5	73.3	115.7
	Optimized	1.8	1.3	109.3	73.7	105.7
FaceQAN	Baseline	1.5	1.9	112.6	72.1	134.6
	Optimized	1.2	1.8	82.8	73.9	98.5
SER-FIQ	Baseline	2.9	2.7	100.0	73.0	/
	Optimized	2.4	2.0	102.2	66.6	/
PCNet	Baseline	1.7	2.3	88.5	59.3	139.7
	Optimized	0.9	1.7	87.6	60.2	116.4
MagFace	Baseline	2.0	3.7	118.6	74.4	294.9
	Optimized	1.7	1.7	125.3	72.3	127.6
SDD-FIQA	Baseline	1.8	2.6	91.4	87.5	190.1
	Optimized	1.8	1.7	80.6	66.6	135.3

(FNMR) at a predefined False Match Rate (FMR), typically fixed at 0.001, at various low-quality image drop (also unconsidered) rates. Specifically, we report the Area Under the ERC Curves (AUC) as our main performance indicator, where smaller values indicate better performance.

Implementation Details. When clustering the embedding space of each individual within the VGGFace2 database, we decide to set the number of clusters C to 20. Consequently, we generate $C - 1 = 19$, mated image pairs for each image, which means that each individual list of mated pairs consists of approximately 60 million pairs. For the hyperparameter θ we use a relatively small value of 0.001, since the goal is to optimize the already computed baseline quality scores. We repeat the whole process 10 times and average the final results.

B. Results

Before presenting results, we note that SER-FIQ was used in the construction of the XQFW database, so any results that combine the two are excluded from the presented analysis.

Same-Model Results. Table I shows the AUC values produced directly with the original FIQA methods (labeled *Baseline*) as well as the AUC scores of the quality-regression network trained using our optimized labels (marked *Optimized*). For readability purposes, the AUC scores are multiplied by 10^3 and rounded to one decimal place. We observe that in most cases the results of our approach are better than those of the underlying FIQA approaches. The only exception to this observation is CR-FIQA, where a concrete improvement is observed only for the hardest of the considered datasets, i.e., XQFW, while the results for the remaining datasets are mostly close, but deteriorate drastically for CPLFW. For all other methods the results consistently improve, with occasional outliers on the CALFW or CPLFW benchmarks.

Cross-Model Results. Table II again shows the AUC values of both the baseline and our (optimized) regression-based FIQA techniques, but this time computed for the cross-model scenario, where the FR model used for estimating the quality of the input images differs from the FR model used for performance reporting. Looking at the individual

TABLE II

CROSS-MODEL PERFORMANCE COMPARISON. COMPARISON OF AUC SCORES ($\text{AUC@FMR1E-3}[\times 10^{-3}](\downarrow)$) BETWEEN THE BASELINE FIQA METHODS (*Baseline*) AND THE PROPOSED OPTIMIZATION APPROACH (*Optimized*), USING THE ARCFACE MODEL FOR QUALITY ESTIMATION AND ELASTICFACE AND CURRICULARFACE FOR PERFORMANCE SCORING. BEST RESULTS ARE HIGHLIGHTED IN **GREEN**.

Methods		ElasticFace					CurricularFace				
		LFW	CFP-FP	CPLFW	CALFW	XQLFW	LFW	CFP-FP	CPLFW	CALFW	XQLFW
CR-FIQA	Baseline	2.1	2.1	62.0	74.9	315.8	1.8	1.9	57.3	76.6	281.0
	Optimized	2.3	2.0	62.7	76.1	305.5	2.0	1.9	56.1	78.2	272.2
FaceQAN	Baseline	1.9	2.6	66.4	77.7	329.6	1.6	2.2	70.7	82.6	323.3
	Optimized	2.0	2.7	65.0	79.7	287.1	1.5	2.3	63.7	83.7	263.5
SER-FIQ	Baseline	3.8	3.9	79.4	74.9	/	3.2	3.4	69.6	78.1	/
	Optimized	3.2	3.0	77.5	69.7	/	2.5	2.5	69.2	71.8	/
PCNet	Baseline	3.2	3.5	74.4	62.0	312.2	2.1	3.0	66.5	64.3	300.3
	Optimized	2.0	2.5	69.5	62.7	295.2	1.3	2.3	62.0	64.7	272.6
MagFace	Baseline	2.6	4.9	79.5	74.9	601.8	2.1	4.4	81.5	79.9	593.2
	Optimized	2.1	2.5	68.6	73.9	306.6	2.0	2.2	65.1	77.6	260.4
SDD-FIQA	Baseline	3.1	3.8	73.1	79.9	480.8	2.1	3.3	64.8	77.4	438.6
	Optimized	2.5	2.5	63.4	68.7	292.8	2.0	2.1	55.7	70.4	268.3




Fig. 4. **Qualitative analysis of the proposed approach.** For each FIQA method, we show the prior distribution of the quality scores of the VGGFace2 database, and an associated scatter plot showing the changes in the quality scores due to our optimization approach.

methods, CR-FIQA and FaceQAN do not show a clear edge for either the baseline or optimized results. While for the hardest benchmark, XQLFW, the optimized variant always performs better than the baseline variant, the opposite is true for CALFW, which contains cross-age image data. For all other FIQA approaches, the proposed optimization method yields better results, and outperforms the baselines in all cases except for PCNet on CALFW. The results are consistent for both the ElasticFace and CurricularFace model.

Cross-Model vs. Same-Model Results. Comparing the cross-model with the same-model results, many similarities can be observed. The performance benefit due to the optimization approach is relatively unconvincing for CR-FIQA, while the results for all other methods are mirrored between the two evaluation schemes. The biggest difference is seen for FaceQAN, where the proposed method performs comparably worse in the cross-model evaluation setting.

Qualitative Analysis. If we look more closely at how the proposed approach works, we see that the distribution of the initial quality scores remains the same under the optimization scheme. This is because the method only rearranges the order of the images and assigns them quality scores from the prior distribution. However, a potential problem with this approach, is that the quality scores of images in higher density areas of the distribution, are harder to change than the quality scores of the images in the lower density areas. This phenomenon is well illustrated in Fig. 4, where for each of the FIQA methods used, a histogram of the prior quality scores over VGGFace2 is presented together with a scatter plot, where each point represents the prior quality of a given image on the x -axis and the optimized quality score on the y -axis. Note how the quality

scores in areas of lower density seem to change drastically, while almost no movement is observed in higher density areas.

Ablation Study. To demonstrate how the optimization of the quality labels affects the final results, we present in Table III AUC scores obtained with a quality-regression network trained with the initial (unoptimized) quality labels as well as the performance gain(-)/loss(+) due to the optimization procedure (in brackets). We use the two most difficult benchmarks: CPLFW and XQLFW, as well as the LFW benchmark for this ablation study. From the presented results, we see that the effectiveness of the optimization in the *same-model scenario*, i.e. with ArcFace, to a certain extent depends on the chosen FIQA technique. For CR-FIQA and SER-FIQ the results do not really seem to favour the optimization approach, as most of the performance gains observed in Table I appear to be a consequence of the transfer learning step. On the *cross-model* side, the results for both ElasticFace and CurricularFace seem to be more in favour of the optimized labels, with only a few counterexamples on the LFW database.

Run-time performance. Because we use a regression-based model trained with the optimized quality scores, the run-time performance of our approach is (approximately) the same regardless of the initial FIQA method used as the basis for the reference quality scores. Thus, the proposed transfer learning step can also be seen as a knowledge distillation procedure that allows us to retain the performance of a given FIQA technique while ensuring a (approximately) fixed run-time complexity, as evidenced by the run-times in Table IV - computed on a desktop PC with an Intel i9-10900KF (3.70GHz) CPU and a Nvidia 3090 GPU with 24GB of video RAM.

TABLE III

RESULTS OF THE ABLATION STUDY. SHOWN ARE AUC SCORES ($\text{AUC@FMR1E-3}[\times 10^{-3}](\downarrow)$) GENERATED BY THE REGRESSION-BASED FIQA METHODS TRAINED WITH THE INITIAL/PRIOR QUALITY LABELS AND THE PERFORMANCE GAIN(-)/LOSS(+) IN TERMS OF AUC CHANGE DUE TO THE LABEL OPTIMIZATION PROCEDURE (IN BRACKETS). RESULTS ARE PRESENTED FOR THE SAME-MODEL (ARCFACE) AND CROSS-MODEL (ELASTICFACE AND CURRICULARFACE) SETTINGS. GAINS ARE MARKED GREEN, LOSSES RED.

Methods	ArcFace			ElasticFace			CurricularFace		
	LFW	CPLFW	XQFW	LFW	CPLFW	XQFW	LFW	CPLFW	XQFW
CR-FIQA	1.3 (+0.5)	84.0 (+1.7)	102.8 (+2.9)	2.1 (+0.2)	65.0 (-2.3)	280.4 (+24.9)	1.6 (+0.4)	64.2 (-4.1)	273.3 (-1.1)
FaceQAN	1.4 (-0.2)	85.7 (-2.9)	109.8 (-11.3)	2.1 (-0.1)	65.5 (-0.5)	298.2 (-11.1)	1.7 (-0.2)	65.2 (-1.5)	273.0 (-9.5)
SER-FIQ	2.2 (+0.2)	101.9 (+0.3)	/	3.2 (-0.0)	79.5 (-2.0)	/	2.4 (+0.1)	72.8 (-3.6)	/
PCNet	0.9 (-0.0)	93.8 (-6.2)	211.3 (-94.9)	2.5 (-0.5)	78.6 (-9.1)	406.6 (-111.4)	1.2 (+0.1)	69.2 (-7.2)	360.1 (-87.5)
MagFace	1.8 (-0.1)	117.2 (+8.1)	195.5 (-67.9)	2.0 (+0.1)	69.8 (-1.2)	492.3 (-185.7)	2.1 (-0.1)	66.1 (-1.0)	501.8 (-241.4)
SDD-FIQA	2.2 (-0.4)	80.9 (-0.3)	164.8 (-29.5)	2.9 (-0.4)	64.3 (-0.9)	393.8 (-101.0)	2.4 (-0.4)	57.5 (-1.8)	337.1 (-68.8)

TABLE IV
RUN-TIME PERFORMANCE. SHOWN IN $[\mu\text{S}]$.

Method	CR-FIQA	FaceQAN	SER-FIQ	PCNet	MagFace	SDD-FIQA
Original	11.5 \pm 5.0	346.5 \pm 9.0	125.2 \pm 19.5	14.5 \pm 1.7	11.5 \pm 4.9	5.7 \pm 5.0
Ours	11.4 \pm 3.0	11.3 \pm 5.0	11.3 \pm 5.0	10.6 \pm 4.9	11.3 \pm 5.0	11.3 \pm 5.0
Speed-up [†]	+0.8%	+3000%	+1100%	+36%	+1.7%	-50%

[†] Approximate values.

V. CONCLUSION

We presented a novel optimization approach, that aims to improve the performance of modern FIQA approaches. A thorough evaluation was performed using multiple state-of-the-art FIQA methods, datasets and FR models. The results of the evaluation showed significant performance improvements in most cases when using the optimization scheme both in the same-model and cross-model setting. As part of our future work, we plan to incorporate multiple sources of quality scores into the optimization procedure to benefit from the complementary quality description provided by different FIQA techniques.

REFERENCES

- [1] M. Wang and W. Deng, "Deep Face Recognition: A Survey," *Neuro-computing*, vol. 429, pp. 215–244, 2021.
- [2] P. Li, L. Prieto, D. Mery, and P. J. Flynn, "On Low-Resolution Face Recognition in the Wild: Comparisons and New Techniques," *IEEE Transactions on Information Forensics and Security*, vol. 14, no. 8, 2019.
- [3] U. Sara, M. Akter, and M. S. Uddin, "Image Quality Assessment Through FSIM, SSIM, MSE and PSNR—a Comparative Study," *Journal of Computer and Communications*, vol. 7, no. 3, pp. 8–18, 2019.
- [4] W. Zhang, K. Ma, J. Yan, D. Deng, and Z. Wang, "Blind Image Quality Assessment Using a Deep Bilinear Convolutional Neural Network," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 1, pp. 36–47, 2018.
- [5] W. Zhang, K. Ma, G. Zhai, and X. Yang, "Learning to Blindly Assess Image Quality in the Laboratory and Wild," in *IEEE International Conference on Image Processing (ICIP)*, pp. 111–115, IEEE, 2020.
- [6] ISO/IEC JTC 1/SC 37 Biometrics, "Information Technology - Biometric Sample Quality - Part 1: Framework," Standard ISO/IEC 29794-1:2016, International Organization for Standardization, 2016.
- [7] T. Schlett, C. Rathgeb, O. Henniger, J. Galbally, J. Fierrez, and C. Busch, "Face Image Quality Assessment: A Literature Survey," *ACM Computing Surveys*, 2022.
- [8] J. Hernandez-Ortega, J. Galbally, J. Fierrez, R. Haraksim, and L. Beslay, "FaceQnet: Quality Assessment for Face Recognition Based on Deep Learning," in *International Conference on Biometrics (ICB)*, 2019.
- [9] O. Fu-Zhao, X. Chen, R. Zhang, Y. Huang, S. Li, J. Li, Y. Li, L. Cao, and W. Yuan-Gen, "SDD-FIQA: Unsupervised Face Image Quality Assessment with Similarity Distribution Distance," in *Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [10] K. Chen, T. Yi, and Q. Lv, "LightQNet: Lightweight Deep Face Quality Assessment for Risk-Controlled Face Recognition," *IEEE Signal Processing Letters*, vol. 28, 2021.
- [11] W. Xie, J. Byrne, and A. Zisserman, "Inducing Predictive Uncertainty Estimation for Face Verification," in *British Machine Vision Conference (BMVC)*, 2020.
- [12] P. Terhörst, J. N. Kolf, N. Damer, F. Kirchbuchner, and A. Kuijper, "SER-FIQ: Unsupervised Estimation of Face Image Quality Based on Stochastic Embedding Robustness," in *Computer Vision and Pattern Recognition (CVPR)*, pp. 5651–5660, 2020.
- [13] Ž. Babnik, P. Peer, and V. Štruc, "FaceQAN: Face Image Quality Assessment Through Adversarial Noise Exploration," in *International Conference on Pattern Recognition (ICPR)*, pp. 748–754, 2022.
- [14] Y. Shi and A. K. Jain, "Probabilistic Face Embeddings," in *International Conference on Computer Vision (ICCV)*, pp. 6902–6911, 2019.
- [15] Q. Meng, S. Zhao, Z. Huang, and F. Zhou, "Magface: A Universal Representation for Face Recognition and Quality Assessment," in *Computer Vision and Pattern Recognition (CVPR)*, pp. 14225–14234, 2021.
- [16] J. G. Cavazos, P. J. Phillips, C. D. Castillo, and A. J. O'Toole, "Accuracy Comparison Across Face Recognition Algorithms: Where are we on Measuring Race Bias?," *IEEE Transactions on Biometrics, Behavior, and Identity Science*, vol. 3, no. 1, pp. 101–111, 2020.
- [17] J. P. Robinson, G. Livitz, Y. Henon, C. Qin, Y. Fu, and S. Timoner, "Face Recognition: Too Bias, or Not Too Bias?," in *Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020.
- [18] X. Gao, S. Z. Li, R. Liu, and P. Zhang, "Standardization of Face Image Sample Quality," in *International Conference on Biometrics (ICB)*, 2007.
- [19] L. Zhang, L. Zhang, and L. Li, "Illumination Quality Assessment for Face Images: A Benchmark and a Convolutional Neural Networks Based Model," in *International Conference on Neural Information Processing (ICONIP)*, pp. 583–593, Springer, 2017.
- [20] Z. Lijun, S. Xiaohu, Y. Fei, D. Pingling, Z. Xiangdong, and S. Yu, "Multi-Branch Face Quality Assessment for Face Recognition," in *International Conference on Communication Technology (ICCT)*, 2019.
- [21] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "VGGFace2: A Dataset for Recognising Faces Across Pose and Age," in *Automatic Face & Gesture Recognition (FG)*, pp. 67–74, 2018.
- [22] F. Boutros, M. Fang, M. Klemm, B. Fu, and N. Damer, "Cr-fiqa: face image quality assessment by learning sample relative classifiability," *arXiv preprint arXiv:2112.06592*, 2021.
- [23] M. Knoche, S. Hoermann, and G. Rigoll, "Cross-Quality LFW: A Database for Analyzing Cross-Resolution Image Face Recognition in Unconstrained Environments," in *IEEE International Conference on Automatic Face and Gesture Recognition (FG 2021)*, pp. 1–5, 2021.
- [24] T. Zheng and W. Deng, "Cross-Pose LFW: A Database for Studying Cross-Pose Face Recognition in Unconstrained Environments," Tech. Rep. 18-01, Beijing University of Posts and Telecommunications, 2018.
- [25] 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 *Winter Conference on Applications of Computer Vision (WACV)*, 2016.
- [26] T. Zheng, W. Deng, and J. Hu, "Cross-Age LFW: A Database for Studying Cross-Age Face Recognition in Unconstrained Environments," *CoRR*, vol. abs/1708.08197, 2017.
- [27] 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," Tech. Rep. 07-49, University of Massachusetts, Amherst, October 2007.
- [28] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "Arcface: Additive Angular Margin Loss for Deep Face Recognition," in *Computer Vision and Pattern Recognition (CVPR)*, pp. 4690–4699, 2019.
- [29] 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 *Computer Vision and Pattern Recognition (CVPR)*, pp. 5901–5910, 2020.
- [30] F. Boutros, N. Damer, F. Kirchbuchner, and A. Kuijper, "ElasticFace: Elastic Margin Loss for Deep Face Recognition," in *Computer Vision and Pattern Recognition (CVPR) Workshops*, pp. 1578–1587, 2022.