

DSL-FIQA: Assessing Facial Image Quality via Dual-Set Degradation Learning and Landmark-Guided Transformer

Wei-Ting Chen^{1,2†} Gurunandan Krishnan² Qiang Gao² Sy-Yen Kuo¹ Sizhuo Ma^{2*} Jian Wang^{2*♦}

¹National Taiwan University ²Snap Inc.

Abstract

Generic Face Image Quality Assessment (GFIQA) evaluates the perceptual quality of facial images, which is crucial in improving image restoration algorithms and selecting high-quality face images for downstream tasks. We present a novel transformer-based method for GFIQA, which is aided by two unique mechanisms. First, a “Dual-Set Degradation Representation Learning” (DSL) mechanism uses facial images with both synthetic and real degradations to decouple degradation from content, ensuring generalizability to real-world scenarios. This self-supervised method learns degradation features on a global scale, providing a robust alternative to conventional methods that use local patch information in degradation learning. Second, our transformer leverages facial landmarks to emphasize visually salient parts of a face image in evaluating its perceptual quality. We also introduce a balanced and diverse Comprehensive Generic Face IQA (CGFIQA-40k) dataset of 40K images carefully designed to overcome the biases, in particular the imbalances in skin tone and gender representation, in existing datasets. Extensive analysis and evaluation demonstrate the robustness of our method, marking a significant improvement over prior methods.

1. Introduction

In the digital era, face images hold a central role in our visual experiences, necessitating a robust metric for assessing their perceptual quality. This metric is crucial for not only evaluating and improving the performance of face restoration algorithms but also for assuring the quality of training datasets for generative models [24, 51]. Designing an effective metric for face image quality assessment presents significant challenges. The inherent complexity of human




Figure 1. **PLCC vs. SRCC Comparison on CGFIQA-40k and GFIQA-20k [59] datasets.** DSL-FIQA, denoted by red triangular points, outperforms other methods (ReIQA [52], StyleGAN-IQA [59], MANIQA [66], TRIQ [63]) and can provide a superior image quality assessment of facial images.

faces, characterized by nuanced visual features and expressions, greatly impacts perceived quality [59]. Additionally, obtaining subjective scores such as Mean Opinion Scores (MOS) is difficult due to the limited availability of licensed face images and the inherent ambiguity in subjective evaluations. Compounding these challenges are facial occlusions caused by masks and accessories, which add another layer of complexity to the assessment process.

Decades of research on image quality assessment (IQA) on general images [14, 28, 35, 66, 71], or general IQA (GIQA), has demonstrated reliable performance across various generic IQA datasets [15, 19, 33, 47, 57, 67]. However, when such methods are applied to faces, they often overlook the distinct features and subtleties inherent to faces, making them less effective for face images.

Another thread of research focuses on biometric face quality assessment (BFIQA) [1, 3, 18, 32, 44, 50, 54, 65], where the goal is to ensure the quality of a given face image for robust biometric recognition. While recognizability is achieved by including factors unique to faces like clarity, pose, and lighting, it does not guarantee accurate assessment of *perceptual* degradation.

A significant stride forward is made by [59], which

Project Page: <https://dsl-fiqa.github.io>

† Part of the work done during internship at Snap Research.

* Co-corresponding authors

♦ Project lead

clearly defines the problem of generic face IQA (GFIQA): GFIQA focuses exclusively on the perceptual quality of face images, as opposed to BFIQA. Their approach leverages pre-trained generative models (StyleGAN2 [26]) to extract latent codes from input images, which are then used as references for quality assessment. Although their method shows promising prediction performance, its effectiveness reduces when input images deviate significantly in shooting angles [59] or quality [46] from the StyleGAN2 training data, limiting its applicability and accuracy to real-world scenarios.

In this paper, we tackle the challenge of GFIQA by proposing a transformer-based method specifically designed to address the limitations of the aforementioned methods. Inspired by modern GIQA techniques [52], we propose a degradation extraction module that obtains degradation representations from input images as intermediate features to aid the regression of quality scores, which is pre-trained via self-supervised learning. However, the existing degradation representation learning scheme [52] often makes an oversimplified assumption that the degradation is uniform across different patches of an image while being distinct from those of other images. This assumption does not hold for real-world data, where diverse degradations within a single image exist due to variations in lighting, motion, camera focus and so on. These inconsistencies may impair the effectiveness of degradation extraction and subsequently hinder the accuracy of quality score prediction.

To this end, we introduce a strategy termed “Dual-Set Degradation Representation Learning” (DSL), which breaks the limits of traditional patch-based learning and extracts degradation representations from a global perspective in degradation learning. This approach is enabled by establishing correspondences between a *controlled* dataset of face images with *synthetic* degradations and a comprehensive *in-the-wild* dataset with *realistic* degradations, offering a comprehensive framework for degradation learning. This degradation representation is injected into the transformer decoder [11] via cross-attention [49], enhancing the overall sensitivity to various kinds of challenging real-world image degradations.

Furthermore, inspired by [64], we utilize the strong correlation between facial image quality and salient facial components such as mouth and eyes. We incorporate landmark detection to localize and feed them as input to our model. This extra module allows our model to autonomously learn to focus on these critical facial components and understand their correlation with the perceptual quality of faces, which helps predict a regional confidence map that aggregates local quality evaluations across the entire face.

Existing datasets such as GFIQA-20k [59] and PIQ23 [5] suffer from limited size or unbalanced distribution. To

bridge this gap, we introduce the Comprehensive Generic Face IQA Dataset (CGFIQA-40k), which comprises 40K images with more balanced diversity in *gender* and *skin tone*. We also include face images with facial occlusions. We believe this dataset will be a valuable resource to fuel and inspire future research.

To summarize, our contributions are as follows:

- We design a transformer-based method specifically designed for GFIQA, predicting perceptual scores for face images.
- We propose “Dual-set Degradation Representation Learning” (DSL), a self-supervised approach for learning degradation features globally. This method effectively captures global degradation representations from both synthetically and naturally degraded images, enhancing the learning process of degradation characteristics.
- We enhance our model’s attention to salient facial components by integrating facial landmark detection, enabling a holistic quality evaluation that adaptively aggregates local quality assessment across the face.
- We construct the Comprehensive Generic Face IQA Dataset (CGFIQA-40k), a collection of 40K face images designed to offer comprehensive and balanced representation in terms of gender and skin tone.

2. Related Work

2.1. Quality Assessment of Face Images

Recent work on Face Image Quality Assessment (FIQA) can be categorized into two major branches [64]: BFIQA and GFIQA. BFIQA originates from biometric studies, focusing on the quality of face images for recognition systems. On the other hand, GFIQA encompasses a wider scope, concentrating on the perceptual degradation of image quality.

Biometric Face Image Quality Assessment (BFIQA): BFIQA evaluates the quality of face images for biometric applications such as face recognition, which often assess images based on established standards [17, 53, 56]. Recent progress in the field has been around learning-based strategies, assessing quality via performance of a recognition system [2, 3, 6, 37, 62]. Some studies have adopted manual labeling, using a set of predefined characteristics as binary constraints [74], while others have investigated subjective aspects of image quality [1]. However, adopting BFIQA methods does not give the best performance when the emphasis is on perceptual quality, which will be demonstrated in the results section.

Generic Face Image Quality Assessment (GFIQA): GFIQA is a recently defined task [59, 64], which prioritizes perceptual degradation in face images instead. Initiatives like Chahine *et al.* [5] highlight the relevance of so-

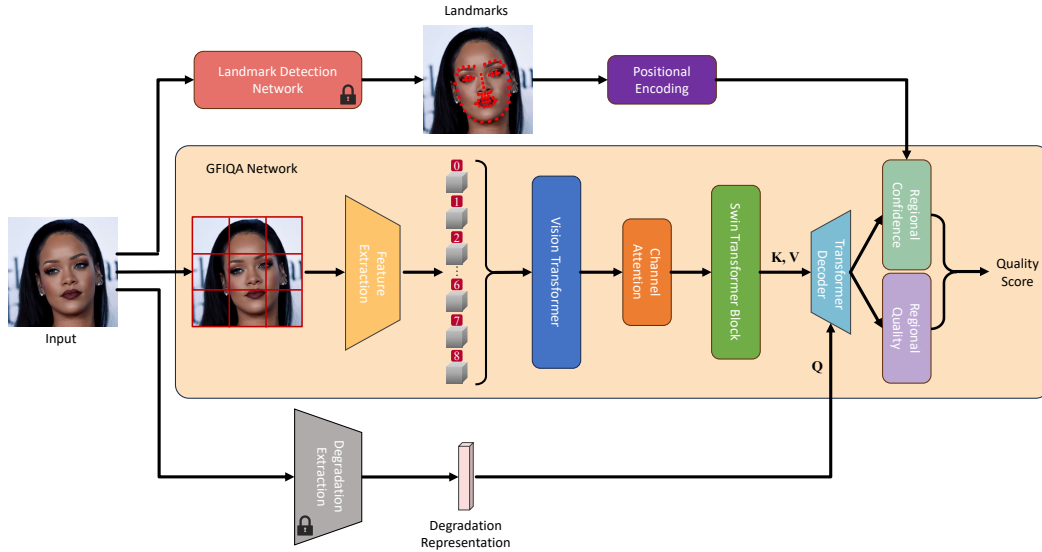


Figure 2. **Overview of our proposed model.** The model contains a core GFIQA network, a degradation extraction network, and a landmark detection network. In our approach, face images are cropped into several patches to fit the input size requirements of the pre-trained ViT feature extractor (See Sec. 3.1). Each patch is then processed individually, and their Mean Opinion Scores (MOS) are averaged to determine the final quality score. For clarity in the figure, the segmentation of the image into patches is not shown.

cial media-driven portrait photography, though their PIQ23 dataset remains restricted in size. Su *et al.* [59] bridges this gap with the larger GFIQA-20k dataset, but it falls short in diversity, lacking in long-tail samples and balanced representation. Their generative prior-based method, while effective, struggles with non-standard images, showing limitations of StyleGAN2-dependent models [25]. This demonstrates the necessity for a more comprehensive dataset and robust GFIQA solution.

2.2. General Image Quality Assessment

Traditional General Image Quality Assessment (GIQA) [7, 14] methods like BRISQUE [39], NIQE [40], and DIVINE [41] are built upon traditional statistical models, which work decently on datasets with constrained size but have faced limitations with complex real-world images. The advent of deep learning has given rise to groundbreaking GIQA methods. RAPIQUE [63] and DB-CNN [69] set new standards in adaptability and accuracy. A further innovation was seen in transformer-based models, including MUSIQ [29] and MANIQA [66], with significantly improved prediction precision. The domain was expanded by CONTRIQUE’s [35] self-supervised paradigm and Zhang *et al.*’s [71] vision-language multitask learning. Saha *et al.* [52] uniquely integrated low and high-level features in an unsupervised manner, emphasizing perceptually relevant quality features.

However, the applicability of these advancements to face images is debatable, as they often overlook facial features

specifically critical for perceptual quality, suggesting a gap ripe for exploration with face-centric quality assessment models.

3. Proposed Method

3.1. Model Overview

Figure 2 illustrates our method: Given an input image $I \in \mathbb{R}^{H \times W \times 3}$, our model estimates its perceptual quality score. In the following, we briefly summarize its components.

Feature Extraction and Refinement: The image initially undergoes feature extraction [66] via a pre-trained Vision Transformer (ViT) [11], followed by a Channel Attention Block [20] that emphasizes relevant inter-channel dependencies. Subsequently, a Swin Transformer Block [34] refines these features, capturing subtle image details.

Degradation Extraction: In parallel, a dedicated module identifies and isolates perceptual degradations within the image, providing a nuanced representation of image quality degradations.

Feature Integration and Quality Estimation: The degradation features, once extracted, are integrated with the outputs from the Swin Transformer within a transformer decoder. This integration employs cross-attention, a technique inspired by Stable Diffusion [49], to enhance the model’s sensitivity to degradation. The combined features are then directed into two MLP branches. The first branch predicts the *regional confidence*, while the second estimates the *regional quality score*. Finally, these outputs are combined

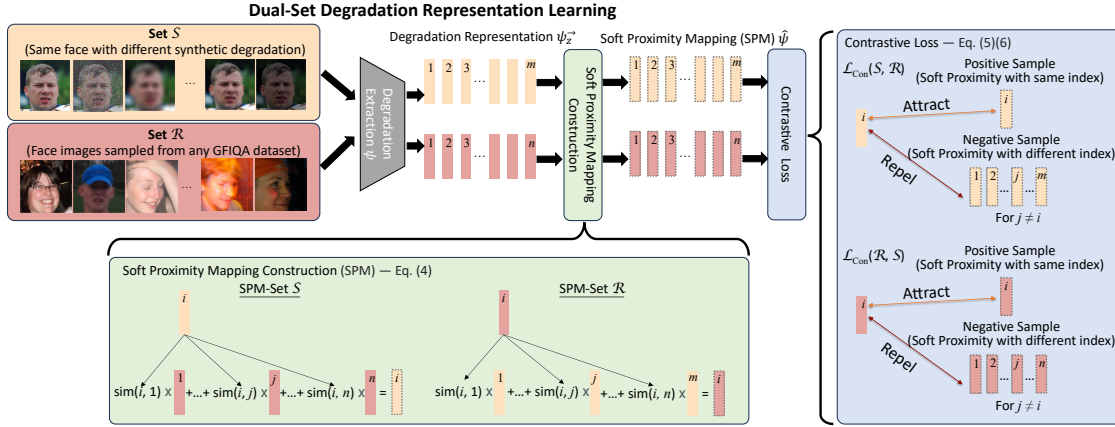


Figure 3. **Dual-Set Degradation Representation Learning (DSL) Illustrated.** On the left, the process of contrastive optimization is depicted, utilizing two unique image sets. Degradation representations are extracted, followed by soft proximity mapping (SPM) calculations and contrastive optimization, compelling the degradation encoder to focus on learning specific degradation features. The right side emphasizes the bidirectional characteristic of our approach, highlighting the comprehensive strategy for identifying and understanding image degradations through contrastive learning.

through a weighted sum to determine the overall quality score of the image.

Landmark-Guided Mechanism: A landmark detection network identifies facial key points, influencing the regional confidence evaluation and ensuring that essential facial features improve the final quality score.

During the training of the core GFIQA network, the landmark and degradation modules remain fixed, leveraging their pre-trained knowledge. Notably, we avoid resizing input images to fit the fixed input dimensions of the pre-trained ViT, which could distort quality predictions [29]. Instead, we crop the image, process each part independently, and average the resulting MOS predictions for a consolidated image quality score. This approach maintains the original dimensions of the image and, consequently, the correctness of perceptual quality assessment.

In the following subsections, we highlight the main technical contribution of our model design: our degradation extraction module and landmark-guided mechanism.

3.2. Self-Supervised Dual-Set Degradation Representation Learning

Before presenting our proposed approach, we provide a concise overview of the existing patch-based degradation learning methods.

3.2.1 Patch-based Degradation Learning

Existing degradation extraction methods ([21, 52, 73]) assume that patches from the same image share similar degradation for contrastive learning. In this framework, patches extracted from the same image are positive samples, while those from different images are negative samples. The

patches are encoded into degradation representations (x , x^+ , and x^-) for the query, positive, and negative samples. The contrastive loss function is designed to enhance the similarity between x and x^+ and dissimilarity between x and x^- , which is given by:

$$\mathcal{L}_{Patch}(x, x^+, x^-) = -\log \frac{\exp(x \cdot x^+ / \theta)}{\sum_{n=1}^N \exp(x \cdot x_n^- / \theta)}, \quad (1)$$

where N is the number of negative samples and θ is a temperature hyper-parameter.

However, the assumption of uniform degradation across the image does not always hold due to lighting, local motion, defocus, and other factors. For example, it is possible to have a moving face with a static background in an image, which means that only some patches suffer from motion blur. This oversimplified assumption often leads to sub-optimal and inconsistent results for degradation learning.

3.2.2 Our Solution

To bridge this gap, we propose Dual-Set Degradation Representation Learning (DSL), which considers entire face images. To make this challenging setting compatible with contrastive learning approaches, we carefully construct two sets of images, \mathcal{S} and \mathcal{R} , each serving a unique purpose in the degradation learning process, as shown in Figure 3.

Set \mathcal{S} consists of a collection of images derived from a single high-quality face image, with each image undergoing different types of Synthetic degradation including but not limited to blurring, noise, resizing, JPEG compression, and extreme lighting conditions. This set acts as a controlled environment, enabling in-depth exploration of a wide variety of degradations against constant content.

In contrast, *Set* \mathcal{R} encompasses a compilation of images from GFIQA datasets, each having *different* content under *Real-world* degradation. This set reflects the unpredictability and diversity of realistic degradations, which are hard to model by synthetic data.

Formally, let $\mathcal{S} = \{s_1, \dots, s_m\}$ and $\mathcal{R} = \{r_1, \dots, r_n\}$, where m and n represent the number of images in \mathcal{S} and \mathcal{R} , respectively. Each image from the two sets is mapped to its degradation representation by a function ψ defined by the degradation extraction module with weights z :

$$\psi_z^{\rightarrow}(\mathcal{S}) = \{\psi(s_1; z), \dots, \psi(s_m; z)\} \quad (2)$$

$$\psi_z^{\rightarrow}(\mathcal{R}) = \{\psi(r_1; z), \dots, \psi(r_n; z)\} \quad (3)$$

Inspired by [48, 58], we introduce a mechanism termed *soft proximity mapping*: For a given image s_i from \mathcal{S} , we map its representation $\psi(s_i)$ to a linear combination of representations in $\psi_z^{\rightarrow}(\mathcal{R})$ as follows:

$$\hat{\psi}(s_i) = \sum_{j=1}^n \text{sim}(\psi(s_i), \psi(r_j)) \cdot \psi(r_j) \quad (4)$$

where $\hat{\psi}(s_i)$ denotes soft proximity mapping of $\psi(s_i)$. $\text{sim}(\cdot, \cdot)$ denotes the similarity between two representations. We use $\mathcal{L}2$ distance as our similarity metric in our implementation. z is omitted for brevity.

This construction allows us to define positive and negative pairs for contrastive learning. Intuitively, a degradation representation $\psi(s_i)$ should be attracted to its own soft proximity mapping $\hat{\psi}(s_i)$, while any other representations $\psi(s_j)$ where $j \neq i$ should be repelled from this soft proximity mapping because s_i and s_j have different degradations by the dedicated construction of *Set* \mathcal{S} . Then, we adopt the contrastive loss [43]:

$$\mathcal{L}_{Con}(\mathcal{S}, \mathcal{R}) = -\frac{1}{m} \sum_{i=1}^m \log \frac{\exp(\text{sim}(\psi(s_i), \hat{\psi}(s_i))/\theta)}{\sum_{j \neq i}^n \exp(\text{sim}(\psi(s_j), \hat{\psi}(s_i))/\theta)} \quad (5)$$

This loss function leverages the nature that within \mathcal{S} , images share the same *content* but differ in *degradations*, contrasting with \mathcal{R} , which varies in both aspects. By drawing the extracted degradation representation closer to its corresponding soft proximity mapping and distancing it from other soft proximity mappings, the degradation extraction module is trained to learn a global degradation representation that is independent of the image content.

Furthermore, the self-supervised **dual-set** contrastive learning strategy is essential for understanding various degradations, particularly in real-world scenarios. This approach is vital as it involves accurately extracting degradation representations from real-world images to approximate those in the synthetic set \mathcal{S} . It is feasible to employ contrastive learning solely on the synthetic set \mathcal{S} to capture degradation patterns: Positive pairs consist of images with

the same degradation, and negative pairs otherwise. However, this naive approach does not generalize well to real-world images. In contrast, our dual-set design can bring together the benefits of both the synthetic set with controllable degradations and the real-world set with realistic degradations, achieving better generalization.

Notice that the roles of \mathcal{S} and \mathcal{R} are symmetric: Just as we utilize representations from \mathcal{S} to seek corresponding features within \mathcal{R} , empirically, we found the reverse is also viable and informative. Thus, we define our Degradation Extraction Loss \mathcal{L}_{DE} as a bidirectional loss:

$$\mathcal{L}_{DE} = \mathcal{L}_{Con}(\mathcal{S}, \mathcal{R}) + \mathcal{L}_{Con}(\mathcal{R}, \mathcal{S}) \quad (6)$$

This bidirectional loss reinforces the mutual learning and alignment between the synthetic and real-world sets, ensuring a comprehensive understanding and representation of realistic degradations. Moreover, it is worth mentioning that the high-quality image in set \mathcal{S} is resampled for every iteration, where this image undergoes random synthetic degradations of varying intensities. Concurrently, images in set \mathcal{R} are also resampled randomly in each iteration.

In summary, DSL gets rid of the uniformity assumption of degradation in patches across the entire image for degradation learning. Instead, it relies on the soft proximity mapping between two constructed sets of images to calculate the contrastive loss, which allows for more precise degradation representation (this mechanism is kind of similar in spirit to [42]). Furthermore, since the entire image is considered, DSL can capture a holistic view of the degradation unique to each image, further boosting the performance.

3.3. Landmark-guided GFIQA

Face images are uniquely challenging in image processing. This is because human eyes are especially sensitive to facial artifacts, raising the importance of nuanced quality assessment [64]. Thus, it is important to design an approach that does not treat each pixel equally; it should acknowledge the perceptual significance of salient facial features. Furthermore, as stated in Sec. 3.1, considering that our network crops the face into various patches to compute the average MOS score, it is crucial to provide landmark information to give the spatial context on which part of the face each patch covers, ensuring a holistic and perceptually consistent evaluation.

As shown in Figure 2, our approach begins with utilizing an existing landmark detection algorithm (*i.e.*, 3DMM model [12]) to identify key facial landmarks. Inspired by Neural Radiance Fields (NeRF) [38], we apply positional encoding to these unique landmark identifiers. By applying a series of sinusoidal functions to the raw identifiers, positional encoding enhances the representational capacity of the network, allowing the network to capture and learn

more intricate relationships and patterns associated with each landmark identifier.

The encoded information is subsequently concatenated with the features processed by the Transformer Decoder, feeding into the regional confidence branch. The human visual system is particularly sensitive to high-frequency details, which are often associated with facial landmarks such as the eyes, nose, and mouth. Providing this landmark-based information to the confidence head can help generate a more precise confidence map, emphasizing regions that humans naturally prioritize in their perception.

In our approach, we deliberately avoid relying on encoding landmark coordinates (x, y) in an image as positions, as it can introduce ambiguity during learning, especially when faces are unaligned, or images are cropped into patches. In such scenarios, specific coordinates may inconsistently correspond to different facial features on different training samples, therefore muddling the learning process. To avoid this, our network employs a fixed encoding scheme for each facial landmark, assigning a unique identifier to every critical feature regardless of its position in the image. This methodology proves particularly advantageous for our ViT, which takes fixed-size crops from the input image, potentially capturing only portions of the face.

Given the diverse range of degradations encountered in GFIQA, off-the-shelf landmark detectors often fail on images with challenging degradations. We observed that fine-tuning existing landmark detectors like [9, 22, 27] on degraded images leads to more accurate landmark detection.

In summary, by adopting landmark-guided cues, our method maintains a consistent awareness of crucial facial features within each crop, which effectively encourages the model to focus on salient facial features when aggregating the regional quality scores.

3.4. Loss Functions

Degradation Encoder. The degradation encoder is trained separately by optimizing Eq. (6). Once trained, it remains fixed when training the core GFIQA network.

GFIQA Network. To measure the discrepancy between the predicted MOS and the ground truth, we employ the Charbonnier loss [31] (\mathcal{L}_{char}), which is defined as:

$$\mathcal{L}_{char}(p, \hat{p}) = \sqrt{(p - \hat{p})^2 + \epsilon^2} \quad (7)$$

where \hat{p} is the predicted MOS, p is the ground truth MOS, and ϵ is a small constant to ensure differentiability.

Unlike existing GIQA [29, 52, 66] or GFIQA [59] models that typically rely on \mathcal{L}_2 losses, we opt for the Charbonnier loss as it is less sensitive to outliers, which in the context of GFIQA can arise from rare face quality degradations, dataset annotation discrepancies, or occasional extreme scores predicted by the model during training. By im-

Dataset	Size	Skin Tone (%)			Gender (%)	
		Light	Medium	Dark	Male	Female
PIQ23	5116	74.5	9.0	16.5	94.0	6.0
GFIQA-20k	20000	81.6	6.7	11.7	64.2	35.8
CGFIQA-40k	39312	53.83	24.91	21.26	51.50	48.50

Table 1. **Comparison of datasets in terms of size, skin tone, and gender distribution.** Skin tones are categorized as Light (Fitzpatrick scale [13] I-II), Medium (Fitzpatrick scale III-IV), and Dark (Fitzpatrick scale V-VI).

proving the robustness against outliers, our model is more aligned with human perceptual judgments.

4. Comprehensive Generic Face IQA Dataset

Existing GFIQA models are evaluated on datasets such as PIQ23 [5] and GFIQA-20k [59]. While PIQ23 contains a variety of in-the-wild images with uncropped faces, its constrained dataset size limits its efficacy for training robust models. Moreover, both datasets exhibit biases in *gender* and *skin tone* representation. This disproportion can introduce biases during model training, decreasing the performance and reliability of models in face image quality assessment tasks. Prior research [4, 8, 30, 55] has shown that this imbalance in data distribution has a significant negative impact on model performance in various face-related applications.

To tackle these challenges, we introduce a new dataset named Comprehensive Generic Face Image Quality Assessment (CGFIQA-40k), which includes approximately 40K images, each with a resolution of 512x512. Each image is annotated by 20 labelers, and each labeler spends about 30 seconds to give a score. From an initial pool of 40,000 images, we filtered out a small number of images with unusable content or incomplete labels, resulting in a total of 39,312 valid images. This dataset is specifically curated to include an extensive collection of face images with diverse distribution on skin tone, gender, and facial occlusions such as masks and accessories.

A comparative overview of our dataset with existing GFIQA datasets is provided in Table 1. We hope this dataset will offer a more comprehensive benchmark for GFIQA, pushing the generalization and robustness of state-of-the-art methods.

5. Experimental Results

5.1. Experiment Settings

Our experiments utilize three datasets: GFIQA-20k [59], PIQ23 [5], and our newly introduced CGFIQA-40k. The GFIQA-20k dataset consists of 20,000 images, divided into 14,000 for training, 2,000 for validation, and 4,000 for testing. PIQ23 contains 5,116 images, with 3,581 for training,

Method	GFIQA-20k		PIQ23		CGFIQA-40k	
	PLCC	SRCC	PLCC	SRCC	PLCC	SRCC
ArcFace [10]	0.9508	0.9510	0.5913	0.6011	0.9722	0.9723
MegaFace [37]	0.9523	0.9531	0.5941	0.5355	0.9731	0.9733
CR-FIQA [3]	0.9593	0.9598	0.6013	0.6021	0.9734	0.9736
Koncept512 [19]	0.9518	0.9523	0.6013	0.6007	0.9713	0.9721
MUSIQ [29]	0.9503	0.9518	0.7141	0.7101	0.9750	0.9735
ReIQA [52]	0.9437	0.9446	0.5988	0.5961	0.9800	0.9802
CONTRIQUE [36]	0.9458	0.9466	0.5892	0.5930	0.9788	0.9799
UNIQUE [70]	0.9413	0.9528	0.5822	0.5710	0.9771	0.9641
MANIQA [66]	0.9614	0.9604	0.7202	0.7180	0.9805	0.9809
TReS [16]	0.9458	0.9552	0.5767	0.5760	0.9816	0.9817
HyperIQA [60]	0.9664	0.9674	0.7152	0.7203	0.9722	0.9733
LIQE [72]	0.9341	0.9571	0.4627	0.4522	0.9786	0.9623
MetaIQA [75]	0.9542	0.9532	0.6008	0.6016	0.9474	0.9463
TRIQ [68]	0.9631	0.9630	0.6118	0.6023	0.9813	0.9821
VCRNet [45]	0.9672	0.9679	0.6826	0.6774	0.9819	0.9821
GraphIQA [61]	0.9434	0.9436	0.4676	0.4709	0.9578	0.9552
IFQA [23]	0.9601	0.9603	0.2907	0.3081	0.9791	0.9803
StyleGAN-IQA [59]	0.9673	0.9684	0.7013	0.7131	0.9822	0.9821
DSL-FIQA	0.9745	0.9740	0.7370	0.7333	0.9873	0.9880

Table 2. Comparison of various image quality assessment methods across three GFIQA datasets. Boldface indicates the best results and underline the second-best. Our approach demonstrates superior and robust performance across all three datasets. Rows in gray denote BFIQA methods, those in pink for generic IQA methods, and in light cyan for GFIQA methods.

Patch	Module							Metric	
	DSL-cat	DSL-S	DSL-R	DSL	Landmark	PE	\mathcal{L}_{Char}	PLCC	SRCC
-	-	-	-	-	-	-	-	0.9682	0.9679
✓	-	-	-	-	-	-	-	0.9695	0.9687
-	✓	-	-	-	-	-	-	0.9703	0.9701
-	-	✓	-	-	-	-	-	0.9713	0.9711
-	-	-	✓	-	-	-	-	0.9709	0.9707
-	-	-	-	✓	-	-	-	0.9721	0.9719
-	-	-	-	✓	✓	-	-	0.9731	0.9728
-	-	-	-	✓	✓	✓	-	0.9735	0.9731
-	-	-	-	✓	✓	✓	✓	0.9745	0.9740

Table 3. Ablation study on proposed techniques. We demonstrate that the proposed modules can effectively improve GFIQA’s performance. “PE” denotes the positional encoding operation. “DSL-S” and “DSL-R” represent Dual-Set Degradation Representation Learning. In this framework, the degradation encoder is trained using a one-sided loss approach, specifically $\mathcal{L}_{Com}(\mathcal{S}, \mathcal{R})$ for “DSL-S” and $\mathcal{L}_{Com}(\mathcal{R}, \mathcal{S})$ for “DSL-R”. “DSL-cat” refers to directly concatenating the degradation representations with image features instead of cross-attention, while “DSL” indicates the adoption of a dual-sided loss function as defined in Eq. (6).

512 for validation, and 1,023 for testing. Our CGFIQA-40k dataset includes a more extensive collection of 39312 images, with a division of 27518 for training, 3931 for validation, and 7863 for testing.

The evaluation metrics employed are the Pearson Linear Correlation Coefficient (PLCC) and Spearman’s Rank Order Correlation Coefficient (SRCC).

5.2. Performance Evaluation

We conducted an extensive evaluation of our method via experiments, comparing it with representative models across the datasets GFIQA-20k, PIQ23, and CGFIQA-40k. To

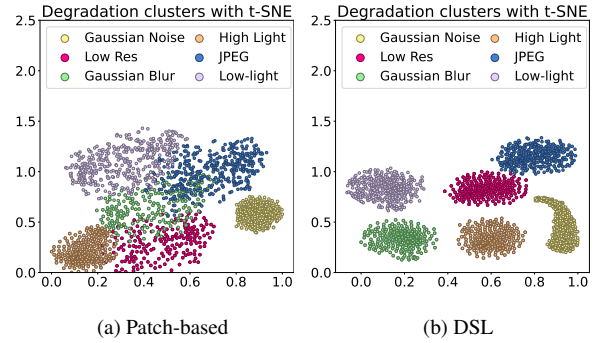


Figure 4. t-SNE visualization of degradation representation extracted using patch-based and DSL-based methods. Unlike the patch-based method, DSL results in well-demarcated clusters for various types of degradation, thereby proving the effectiveness of the learned representations.

Strategy	Patch-based	DSL
mAP	48.2	69.3

Table 4. Comparative retrieval accuracy using DSL and patch-based methods under real-world degradations, quantified by mAP scores.

maintain a fair comparison, all models were trained and validated under the identical conditions specified in Sec. 5.1.

We compare with a wide range of generic IQA models, including Koncept512 [19], MUSIQ [29], ReIQA [52], CONTRIQUE [36], UNIQUE [70], MANIQA [66], TReS [16], HyperIQA [60], LIQE [72], MetaIQA [75], TRIQ [68], VCRNet [45], and GraphIQA [61]. We also compare with recent GFIQA methods, including StyleGAN-IQA [59] and IFQA [23]. For completeness, we include three representative BFIQA methods, ArcFace [10], MegaFace [37], and CR-FIQA [3].

Table 2 clearly shows the robust performance of our method, which outperforms existing models across all metrics on GFIQA-20k, PIQ23, and CGFIQA-40k. The consistent results across diverse datasets validate our approach’s strength and adaptability, establishing it as a robust generic face image quality assessment solution.

5.3. Ablation Study

To validate the effectiveness of individual components of our approach, we conducted an ablation study based on the GFIQA-20k dataset.

Effectiveness of Degradation Extraction. The innovation of our method lies in the Dual-set Degradation Representation Learning (DSL), enhancing GFIQA precision by capturing global degradation in degradation learning. To verify its impact, we show the comparison in Table 3. Including DSL improved the performance of the GFIQA (Row 1

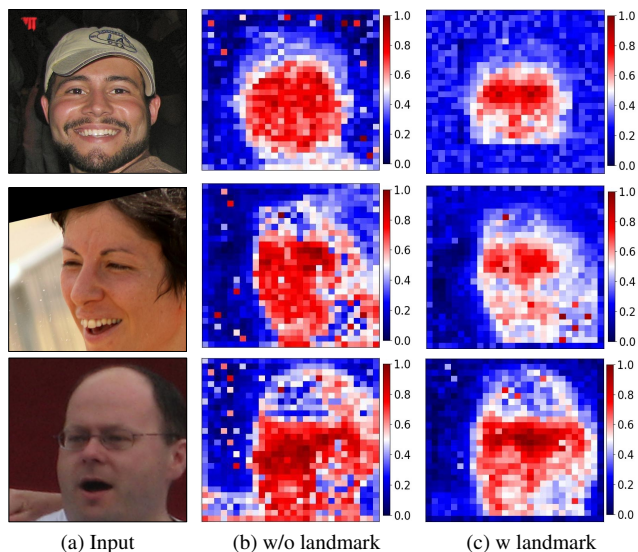


Figure 5. **Comparison of using landmark mechanism to guide the GFIQA network.** We present the regional confidence maps and the corresponding input. With landmark guidance, the confidence maps focus more on key facial landmarks, providing a more discriminative assessment. In contrast, without landmark guidance, the confidence maps tend to cover the entire face, often lacking specificity and even assigning higher confidence to irrelevant areas (e.g., background).

and Row 6), highlighting its critical role in boosting GFIQA accuracy. We also substituted the DSL learning technique for the existing patch-based degradation learning in Table 3 (Row 2 and Row 6), and the results indicate that the proposed DSL can improve GFIQA performance more than the patch-based strategy (Sec. 3.2.1). We also validated the significance of cross-attention for integrating degradation information (Row 3 and Row 6). The results indicated that employing cross-attention for this integration yielded superior outcomes.

Comparison between DSL and patch-based methods. To further evaluate the advantage of our DSL over patch-based degradation extraction methods, we conducted two additional experiments. Both methods were trained on the same dataset for fair and direct comparison. The testing data used was completely independent of the training set, guaranteeing the validity of our assessment.

The first experiment involved synthetic data. We randomly selected 1,000 face images from the FFHQ dataset [25], subjecting each to six types of synthetic degradations. We then employed DSL and patch-based feature extraction, subsequently visualizing these features using t-SNE. The results, illustrated in Figure 4, demonstrated that DSL could effectively separate the images based on their specific degradations, while the patch-based method showed considerable overlap.

Furthermore, we extended our exploration to real-world conditions, using images from the GFIQA-20k dataset. This second experiment was designed to verify if the distinct degradation representations learned by DSL could enhance image degradation retrieval accuracy under real-world degradations. To this end, we synthesized six types of degradations on 100 images from the FFHQ dataset. These synthetically degraded images were used as queries to probe the GFIQA-20k test set, selecting the top 5 images with the smallest distance. We then verify whether the five images fall under the same degradation category by human inspection. We quantified our method’s precision by calculating the mean average precision (mAP) for these retrieval tasks, as shown in Table 4. The results confirmed DSL’s enhanced accuracy in identifying images with similar degradation attributes.

In conclusion, these experiments emphasized the effectiveness of DSL in crafting distinct degradation representations and its practical superiority in real-world scenarios, bolstering its value in improving GFIQA outcomes.

Effectiveness of Landmark-guided GFIQA. Integrating facial landmarks into GFIQA significantly improves quality assessment accuracy, addressing the complexity of facial features often ignored in traditional methods, which is validated in Table 3 (Row 6 and 7). To understand how the landmark guidance works, Figure 5 visualizes the regional confidences predicted with and without landmarks guidance. When landmarks are not used, the model indiscriminately overemphasizes the entire face and even background areas. In contrast, the model with landmark guidance focuses on crucial facial regions, which are more aligned with human perception. In addition, Table 3 (Row 7 and 8) substantiate the benefit of applying positional encoding to the landmark identifiers, showing that positional encoding can indeed enhance the model capacity to capture more complex relationships inherent in facial features, thereby improving the overall prediction accuracy.

Effectiveness of Charbonnier Loss. The introduction of Charbonnier loss improves the accuracy of GFIQA as shown in Table 3 (Row 8 and 9).

6. Conclusion

In this paper, we tackle the inherent complexities in GFIQA with a transformer-based method. Our Dual-Set Degradation Representation Learning improves degradation extraction, and the additional guidance from facial landmarks further improves the assessment accuracy. Furthermore, we curate the CGFIQA-40k Dataset, rectifying imbalances in skin tones and gender ratios prevalent in previous datasets. Extensive experimental results show that the proposed method performs favorably against state-of-the-art methods across several GFIQA datasets.

References

- [1] Lacey Best-Rowden and Anil K Jain. Learning face image quality from human assessments. *IEEE TIFS*, 2018. 1, 2
- [2] Samarth Bharadwaj, Mayank Vatsa, and Richa Singh. Can holistic representations be used for face biometric quality assessment? In *ICIP*, 2013. 2
- [3] Fadi Boutros, Meiling Fang, Marcel Klemt, Biying Fu, and Naser Damer. Cr-fiq: face image quality assessment by learning sample relative classifiability. In *CVPR*, 2023. 1, 2, 7
- [4] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 2018. 6
- [5] Nicolas Chahine, Stefania Calarasanu, Davide Garcia-Civiero, Théo Cayla, Sira Ferradans, and Jean Ponce. An image quality assessment dataset for portraits. In *CVPR*, 2023. 2, 6
- [6] Jiansheng Chen, Yu Deng, Gaocheng Bai, and Guangda Su. Face image quality assessment based on learning to rank. *IEEE Signal Processing Letters*, 2014. 2
- [7] Zewen Chen, Juan Wang, Bing Li, Chunfeng Yuan, Weihua Xiong, Rui Cheng, and Weiming Hu. Teacher-guided learning for blind image quality assessment. In *ACCV*, 2022. 3
- [8] Cynthia M Cook, John J Howard, Yevgeniy B Sirotnin, Jerry L Tipton, and Arun R Vemury. Demographic effects in facial recognition and their dependence on image acquisition: An evaluation of eleven commercial systems. *IEEE Transactions on Biometrics, Behavior, and Identity Science*, 2019. 6
- [9] Jiansheng Chen, Jia Guo, Evangelos Ververas, Irene Kotsia, and Stefanos Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *CVPR*, 2020. 6
- [10] Jiansheng Chen, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019. 7
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 2, 3
- [12] Bernhard Egger, William AP Smith, Ayush Tewari, Stefanie Wuhler, Michael Zollhoefer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, et al. 3d morphable face models—past, present, and future. *ACM ToG*, 2020. 5
- [13] Thomas B Fitzpatrick. Soleil et peau. *J. Med. Esthet.*, 1975. 6
- [14] Yixuan Gao, Xionghuo Min, Wenhan Zhu, Xiao-Ping Zhang, and Guangtao Zhai. Image quality score distribution prediction via alpha stable model. *IEEE TCSVT*, 2022. 1, 3
- [15] Deepti Ghadiyaram and Alan C Bovik. Massive online crowdsourced study of subjective and objective picture quality. *IEEE Transactions on Image Processing*, 25(1):372–387, 2015. 1
- [16] S Alireza Golestaneh, Saba Dadsetan, and Kris M Kitani. No-reference image quality assessment via transformers, relative ranking, and self-consistency. In *WACV*, 2022. 7
- [17] P Grother, Austin Hom, Mei Ngan, and Kayee Hanaoka. Ongoing face recognition vendor test (FRVT) part 5: Face image quality assessment. *NIST Interagency Report*, 2020. 2
- [18] Javier Hernandez-Ortega, Javier Galbally, Julian Fierrez, Rudolf Haraksim, and Laurent Beslay. Faceqnet: Quality assessment for face recognition based on deep learning. In *ICB*. IEEE, 2019. 1
- [19] Vlad Hosu, Hanhe Lin, Tamas Sziranyi, and Dietmar Saupe. Koniq-10k: An ecologically valid database for deep learning of blind image quality assessment. *TIP*, 2020. 1, 7
- [20] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018. 3
- [21] Zhi-Kai Huang, Wei-Ting Chen, Yuan-Chun Chiang, Sy-Yen Kuo, and Ming-Hsuan Yang. Counting crowds in bad weather. In *ICCV*, 2023. 4
- [22] Haibo Jin, Shengcai Liao, and Ling Shao. Pixel-in-pixel net: Towards efficient facial landmark detection in the wild. *IJCV*, 2021. 6
- [23] Byungho Jo, Donghyeon Cho, In Kyu Park, and Sungeun Hong. Ifqa: Interpretable face quality assessment. In *WACV*, 2023. 7
- [24] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 1
- [25] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 3, 8
- [26] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020. 2
- [27] Yury Kartynnik, Artsiom Ablavatski, Ivan Grishchenko, and Matthias Grundmann. Real-time facial surface geometry from monocular video on mobile gpus. *arXiv preprint arXiv:1907.06724*, 2019. 6
- [28] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. MUSIQ: multi-scale image quality transformer. 2021. 1
- [29] Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *ICCV*, 2021. 3, 4, 6, 7
- [30] Brendan F Klare, Mark J Burge, Joshua C Klontz, Richard W Vorder Bruegge, and Anil K Jain. Face recognition performance: Role of demographic information. *TIFS*, 2012. 6
- [31] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *CVPR*, 2017. 6
- [32] Zhang Lijun, Shao Xiaohu, Yang Fei, Deng Pingling, Zhou Xiangdong, and Shi Yu. Multi-branch face quality assessment for face recognition. In *ICCT*. IEEE, 2019. 1
- [33] Hanhe Lin, Vlad Hosu, and Dietmar Saupe. Kadid-10k: A large-scale artificially distorted iqa database. In *QoMEX*. IEEE, 2019. 1
- [34] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 3
- [35] Pavan C. Madhusudana, Neil Birkbeck, Yilin Wang, Balu Adsumilli, and Alan C. Bovik. Image quality assessment

- using contrastive learning. [abs/2110.13266](https://arxiv.org/abs/2110.13266), 2021. 1, 3
- [36] Pavan C Madhusudana, Neil Birkbeck, Yilin Wang, Balu Adsumilli, and Alan C Bovik. Image quality assessment using contrastive learning. *TIP*, 2022. 7
- [37] Qiang Meng, Shichao Zhao, Zhida Huang, and Feng Zhou. Magface: A universal representation for face recognition and quality assessment. In *CVPR*, 2021. 2, 7
- [38] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 2021. 5
- [39] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *TIP*, 2012. 3
- [40] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal processing letters*, 2012. 3
- [41] Anush Krishna Moorthy and Alan Conrad Bovik. Blind image quality assessment: From natural scene statistics to perceptual quality. *TIP*, 2011. 3
- [42] Kieran A Murphy, Varun Jampani, Srikumar Ramalingam, and Ameesh Makadia. Learning abcs: Approximate bijective correspondence for isolating factors of variation with weak supervision. In *CVPR*, 2022. 5
- [43] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 5
- [44] Fu-Zhao Ou, Xingyu Chen, Ruixin Zhang, Yuge Huang, Shaoxin Li, Jilin Li, Yong Li, Liujuan Cao, and Yuan-Gen Wang. Sdd-fqa: unsupervised face image quality assessment with similarity distribution distance. In *CVPR*, 2021. 1
- [45] Zhaoqing Pan, Feng Yuan, Jianjun Lei, Yuming Fang, Xiao Shao, and Sam Kwong. Vcnet: Visual compensation restoration network for no-reference image quality assessment. *TIP*, 2022. 7
- [46] Yohan Poirier-Ginter and Jean-François Lalonde. Robust unsupervised stylegan image restoration. In *CVPR*, 2023. 2
- [47] Nikolay Ponomarenko, Vladimir Lukin, Alexander Zelen-sky, Karen Egiazarian, Marco Carli, and Federica Battisti. Tid2008-a database for evaluation of full-reference visual quality assessment metrics. *Advances of modern radioelec-tronics*, 2009. 1
- [48] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. *NeurIPS*, 2018. 5
- [49] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 3
- [50] Jacob Rose and Thirimachos Bourlai. Deep learning based estimation of facial attributes on challenging mobile phone face datasets. In *IEEE/ACM international conference on advances in social networks analysis and mining*, 2019. 1
- [51] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *CVPR*, 2023. 1
- [52] Avinab Saha, Sandeep Mishra, and Alan C Bovik. Re-iqa: Unsupervised learning for image quality assessment in the wild. In *CVPR*, 2023. 1, 2, 3, 4, 6, 7
- [53] Torsten Schlett, Christian Rathgeb, Olaf Henniger, Javier Galbally, Julian Fierrez, and Christoph Busch. Face image quality assessment: A literature survey. *ACM Computing Surveys*, 2021. 2
- [54] Torsten Schlett, Christian Rathgeb, Olaf Henniger, Javier Galbally, Julian Fierrez, and Christoph Busch. Face image quality assessment: A literature survey. *ACM Computing Surveys (CSUR)*, 2022. 1
- [55] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. Fairness and abstraction in sociotechnical systems. In *Conference on fairness, accountability, and transparency*, 2019. 6
- [56] Harin Sellahewa and Sabah A Jassim. Image-quality-based adaptive face recognition. *TMM*, 2010. 2
- [57] Hamid R Sheikh, Muhammad F Sabir, and Alan C Bovik. A statistical evaluation of recent full reference image quality assessment algorithms. *TIP*, 2006. 1
- [58] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *NeurIPS*, 2017. 5
- [59] Shaolin Su, Hanhe Lin, Vlad Hosu, Oliver Wiedemann, Jinqiu Sun, Yu Zhu, Hantao Liu, Yanning Zhang, and Dietmar Saupe. Going the extra mile in face image quality assessment: A novel database and model. *TMM*, 2023. 1, 2, 3, 6, 7
- [60] Shaolin Su, Qingsen Yan, Yu Zhu, Cheng Zhang, Xin Ge, Jinqiu Sun, and Yanning Zhang. Blindly assess image quality in the wild guided by a self-adaptive hyper network. In *CVPR*, 2020. 7
- [61] Simen Sun, Tao Yu, Jiahua Xu, Wei Zhou, and Zhibo Chen. Graphiqa: Learning distortion graph representations for blind image quality assessment. *TMM*, 2022. 7
- [62] Philipp Terhorst, Jan Niklas Kolf, Naser Damer, Florian Kirchbuchner, and Arjan Kuijper. Ser-fiq: Unsupervised estimation of face image quality based on stochastic embedding robustness. In *CVPR*, 2020. 2
- [63] Zhengzhong Tu, Xiangxu Yu, Yilin Wang, Neil Birkbeck, Balu Adsumilli, and Alan C Bovik. Rapique: Rapid and accurate video quality prediction of user generated content. *IEEE Open Journal of Signal Processing*, 2021. 1, 3
- [64] Tao Wang, Kaihao Zhang, Xuanxi Chen, Wenhan Luo, Jiankang Deng, Tong Lu, Xiaochun Cao, Wei Liu, Hongdong Li, and Stefanos Zafeiriou. A survey of deep face restoration: Denoise, super-resolution, deblur, artifact removal. *arXiv preprint arXiv:2211.02831*, 2022. 2, 5
- [65] Fei Yang, Xiaohu Shao, Lijun Zhang, Pingling Deng, Xi-angdong Zhou, and Yu Shi. Dfqa: Deep face image quality assessment. In *ICIG*, 2019. 1
- [66] Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. Maniqa: Multi-dimension attention network for no-reference image quality assessment. In *CVPR*, 2022. 1, 3, 6, 7
- [67] Zhenqiang Ying, Haoran Niu, Praful Gupta, Dhruv Mahajan, Deepti Ghadiyaram, and Alan Bovik. From patches to pictures (paq-2-piq): Mapping the perceptual space of picture quality. In *CVPR*, 2020. 1
- [68] Junyong You and Jari Korhonen. Transformer for image quality assessment. In *ICIP*, 2021. 7
- [69] Weixia Zhang, Kede Ma, Jia Yan, Dexiang Deng, and Zhou

- Wang. Blind image quality assessment using a deep bilinear convolutional neural network. *TCSVT*, 2018. [3](#)
- [70] Weixia Zhang, Kede Ma, Guangtao Zhai, and Xiaokang Yang. Uncertainty-aware blind image quality assessment in the laboratory and wild. *TIP*, 2021. [7](#)
- [71] Weixia Zhang, Guangtao Zhai, Ying Wei, Xiaokang Yang, and Kede Ma. Blind image quality assessment via vision-language correspondence: A multitask learning perspective. In *CVPR*, 2023. [1, 3](#)
- [72] Weixia Zhang, Guangtao Zhai, Ying Wei, Xiaokang Yang, and Kede Ma. Blind image quality assessment via vision-language correspondence: A multitask learning perspective. In *CVPR*, 2023. [7](#)
- [73] Kai Zhao, Kun Yuan, Ming Sun, Mading Li, and Xing Wen. Quality-aware pre-trained models for blind image quality assessment. In *CVPR*, 2023. [4](#)
- [74] Xuan Zhao, Yali Li, and Shengjin Wang. Face quality assessment via semi-supervised learning. In *ICPR*, 2019. [2](#)
- [75] Hancheng Zhu, Leida Li, Jinjian Wu, Weisheng Dong, and Guangming Shi. Metaiqa: Deep meta-learning for no-reference image quality assessment. In *CVPR*, 2020. [7](#)